An Interop Publication

ConneXions

The Interoperability Report

March 1992 Volume 6, No. 3

ConneXions — The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

ConneXions is published monthly by Interop, Inc., 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. 415-941-3399. Fax: 415-949-1779. Toll-free: 1-800-INTEROP.

Copyright © 1992 by Interop, Inc. Quotation with attribution encouraged.

ConneXions—The Interoperability Report and the *ConneXions* logo are registered trademarks of Interop, Inc.

ISSN 0894–5926

From the Editor

This month marks the 5th anniversary of the first INTEROP conference (called "TCP/IP Interoperability Conference" in those days). Since we only distributed a 12 page "Premiere Issue" at that conference, it isn't time just yet to call this the Anniversary Issue. (That happens in May, and will coincide with INTEROP 92 Spring).

As I prepare for our 7th INTEROP, and the 5th Anniversary Issue, I've been pondering some of the changes that have taken place in the networking world in the last few years. One thing that immediately stands out is a dramatic "paradigm shift" with respect to protocol stacks. In the early days, it was generally accepted that networks would eventually transition to OSI, but would use TCP/IP as the "interim" multi-vendor solution. Today, the picture is not quite so simple. Instead of "TCP/IP or OSI," people are now saying "multi-protocol" and "...make all the pieces work together."

With this in mind, it is only natural that we look at some of the major proprietary networking solutions, and what better way to start than with IBM's *Systems Network Architecture* (SNA). Wayne Clark of Cisco Systems introduces us to SNA Internetworking.

At every INTEROP exhibition we've organized a number of *Solutions Showcase Demonstrations*™ where vendors cooperate to show a particular technology working across the INTEROP shownet (and in some case across the world). We asked one OSI demo participant, Sue Hares, to give us an inside look at what it is like to plan and execute such a complex event.

It is encouraging to see the *Metamail* system become generally available (see page 20). This software, along with extensions to RFC 822, could render X.400 obsolete in the Internet community. Time will tell.

There are many ways to get copies of RFCs over the Internet (see *ConneXions,* Volume 6, No. 1, January 1992). Most of these simply access a directory of files where each RFC is a file. The searching capability (if any) is limited to the filename recognition features of that system. Other systems like *WAIS* (see *ConneXions,* Volume 5, No. 11, November 1991) and *Archie* (see *ConneXions,* Volume 6, No. 2, February 1992) offer somewhat more searching ability, though filenames are still the commonly available search key. Another system for finding the RFC you want has recently been introduced—the *ISI RFC-Info* server. In this system you can search for an RFC by author, date, or keyword (all title words are automatically keywords). More information about this new service can be found on page 22.

SNA Internetworking

by Wayne Clark, Cisco Systems, Inc.

Introduction

Since its introduction by IBM in 1974, *Systems Network Architecture* (SNA) has traditionally been the networking architecture of choice for commercial computing. Today, there are over 40,000 operational SNA networks worldwide. It is estimated that 85% of all Fortune 1000 companies in the United States are based upon SNA [1].

SNA was IBM's grand unification effort to consolidate a variety of mainframe-based access methods and host-to-terminal protocols into a single coherent, flexible architecture. Given its host-centric orientation, SNA has predominantly been a hierarchical networking architecture. The mainframe-based elements of SNA wield the greatest power within the entire network while significantly less autonomy is granted to the remote SNA nodes. In general, the further a node gets away from the host in the SNA network, the fewer SNA capabilities that node contains.

This topology served the needs of commercial computing during the 1970s and early 1980s. Over the past several years, SNA has been evolving to meet the demands of today's highly distributed computing environment. In the mid 1980s, IBM introduced a peer-based form of SNA known as *Advanced Peer-to-Peer Networking* (APPN) on its mid-range computer, the System/36. Though it was originally a prototype architecture, this variant of SNA has been released upon more and more IBM platforms lending credibility to IBM's statement that APPN will eventually replace (or at least coexist with) the original hierarchical SNA.

Internetworking with SNA

When SNA was first introduced, it had vastly greater capabilities than the other networking alternatives of its day. However, SNA really hasn't kept pace with the price/performance achievements afforded other protocols such as IPX, TCP/IP, DECnet [11], and AppleTalk. With the proliferation of local area networks and the success of other networking architectures (especially TCP/IP), internetworking has begun to encroach on the commercial computing environment once exclusively controlled by SNA. This encroachment has resulted in two different overall internetworking structures relative to SNA:

- Support of LAN internetworks beneath SNA, and
- Accommodation of SNA within other WAN architectures.

Both variants permit corporations to retain their SNA end systems and thus preserve their application investment with minimal impact to the users. IBM champions the first form of SNA internetworking and offers a variety of products while third-party vendors are in hot pursuit and, in most cases, surpassing IBM's functionality. Multiprotocol router vendors also offer products that conform to the second structure by tunnelling SNA frames across disparate wide area networks.

While the details of both variants are imperceptible to the end user (with the probable exception of improved performance), the second alternative provides a novel set of benefits and challenges to the SNA network administrator. This article will discuss both variants of SNA internetworking but will give more attention to the second form since matching the needs of SNA to the characteristics of other WAN architectures is a blessing which comes with some risks.

SNA terminology

Those of you who have come into casual contact with SNA realize that the terminology space for SNA is immense. Fortunately, only a small subset of terms is required in order to understand its interoperability with other architectures. That subset is presented below.

Subarea Network: A hierarchical SNA network consisting of interconnected hosts and communications controllers. Subarea network nodes can be either Node Type 5 or Node Type 4.

APPN Network: A peer-to-peer network consisting of interconnected nodes that implement IBM's Node Type 2.1. APPN networks support three different types of nodes.

> *Network Nodes (NN)* are intermediate nodes that perform route selection and provide directory services to other APPN nodes.
>
> *End Nodes (EN)* are end systems that can be the source and/or target, but do not provide any routing services. End Nodes rely on their directly-connected Network Node for APPN services.
>
> *Low Entry Networking Nodes (LEN)* are end systems similar to End Nodes but cannot rely on Network Nodes for APPN services and therefore must have a statically-defined image of the APPN network.

Physical Unit (PU): Representation of a node in an SNA network. A Physical Unit is responsible for managing the communication resources of the node. The SNA capabilities of a node differs depending upon which type of Physical Unit (aka *Node Type*) it embodies.

> *Node Type 5* is the most comprehensive of all node types and is contained within the host access method, *Virtual Telecommunications Access method* (VTAM), on IBM mainframes. Node Type 5 embodies the *System Services Control Point* (SSCP), the SNA component responsible for controlling an entire hierarchical SNA network.
>
> *Node Type 4* resides on IBM's communications controllers are used to route and control the flow of data through an SNA subarea network. The best-known example of an IBM communications controller is the IBM 3745.
>
> *Node Type 2* is the least powerful SNA node type and executes on a communications device on the periphery of the SNA subarea network. The best-known example of a Node Type 2 device is the IBM 3x74 cluster controller used for concentrating IBM terminals and printers.
>
> *Node Type 2.1* is a relatively new type of node used for connecting SNA nodes in a peer-oriented network. APPN is based upon Node Type 2.1. Type 2.1 nodes can also be connected into a traditional hierarchical SNA network.

Figure 1 on the following page shows a sample SNA subarea network while Figure 2 shows a sample APPN network. In Figure 1, the interconnected PU 5 and PU 4 nodes form the SNA subarea network. Also, notice in Figure 2 that the combination of a host node and its communications controller are a Low Entry Networking Node (LEN) to the APPN network.

continued on next page

SNA Internetworking *(continued)*

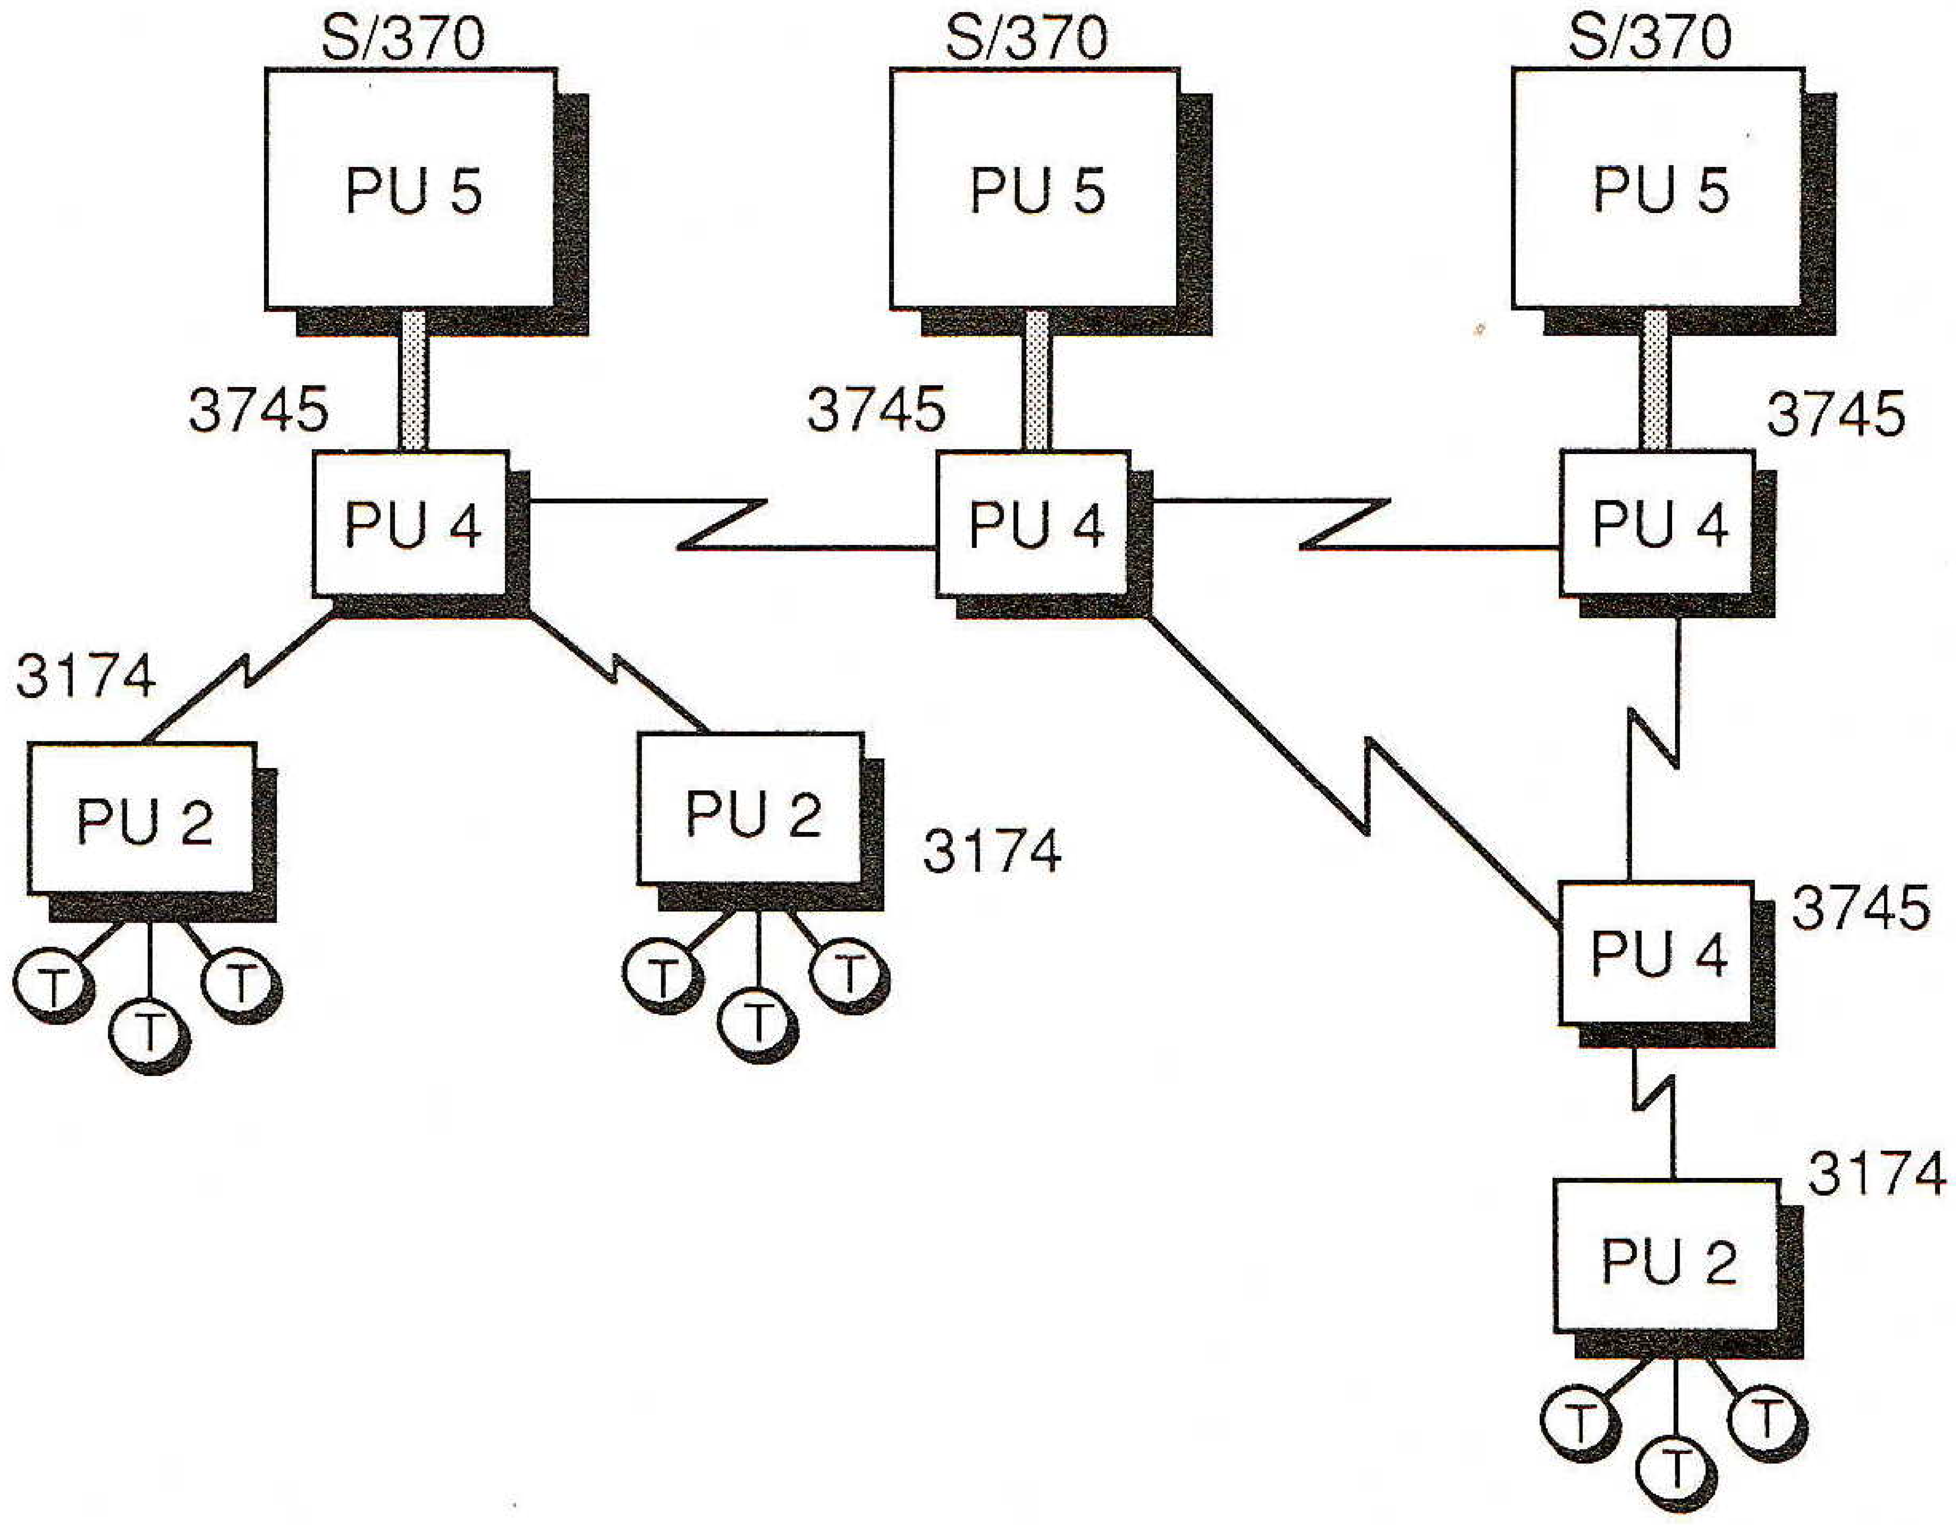

Figure 1: An SNA Subarea Network

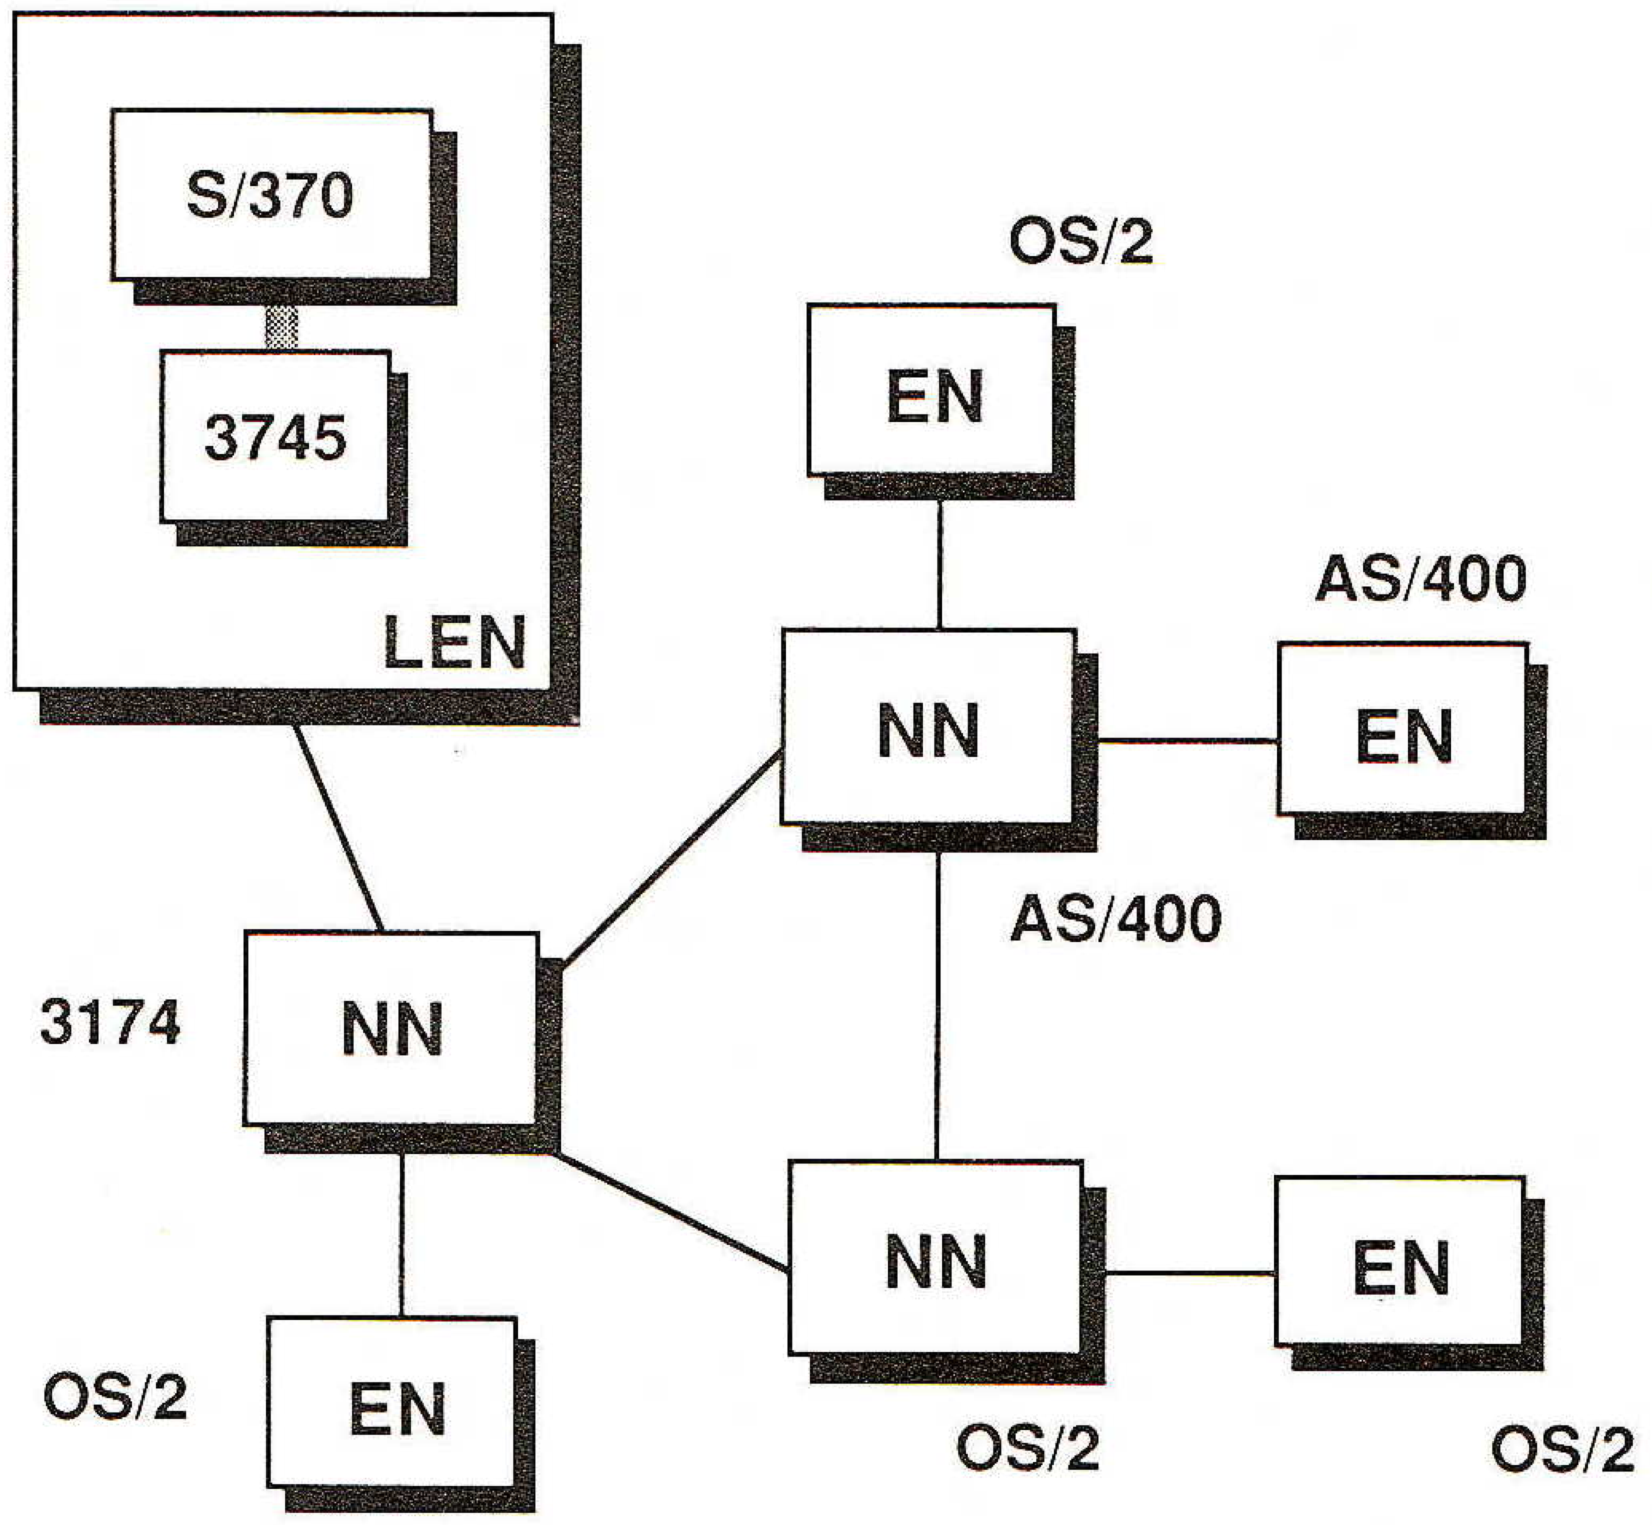

Figure 2: An APPN Network

SNA layers

It should come as no surprise that SNA, like all modern networking protocols, is a layered architecture. In fact, IBM furnished SNA as a working example of a layered architecture to the *International Organization for Standardization* (ISO) in the late 1970s. It should therefore come as no surprise that some of OSI's constructs have their origins in SNA.

Since SNA and OSI are the result of different design teams, there are not only differences in the services provided but also differences in the layer distribution of services that are similar. The layers of SNA and OSI are shown in Figure 3. A detailed comparison of the layers of SNA and OSI can be found in Cypser [3] and Martin [4].

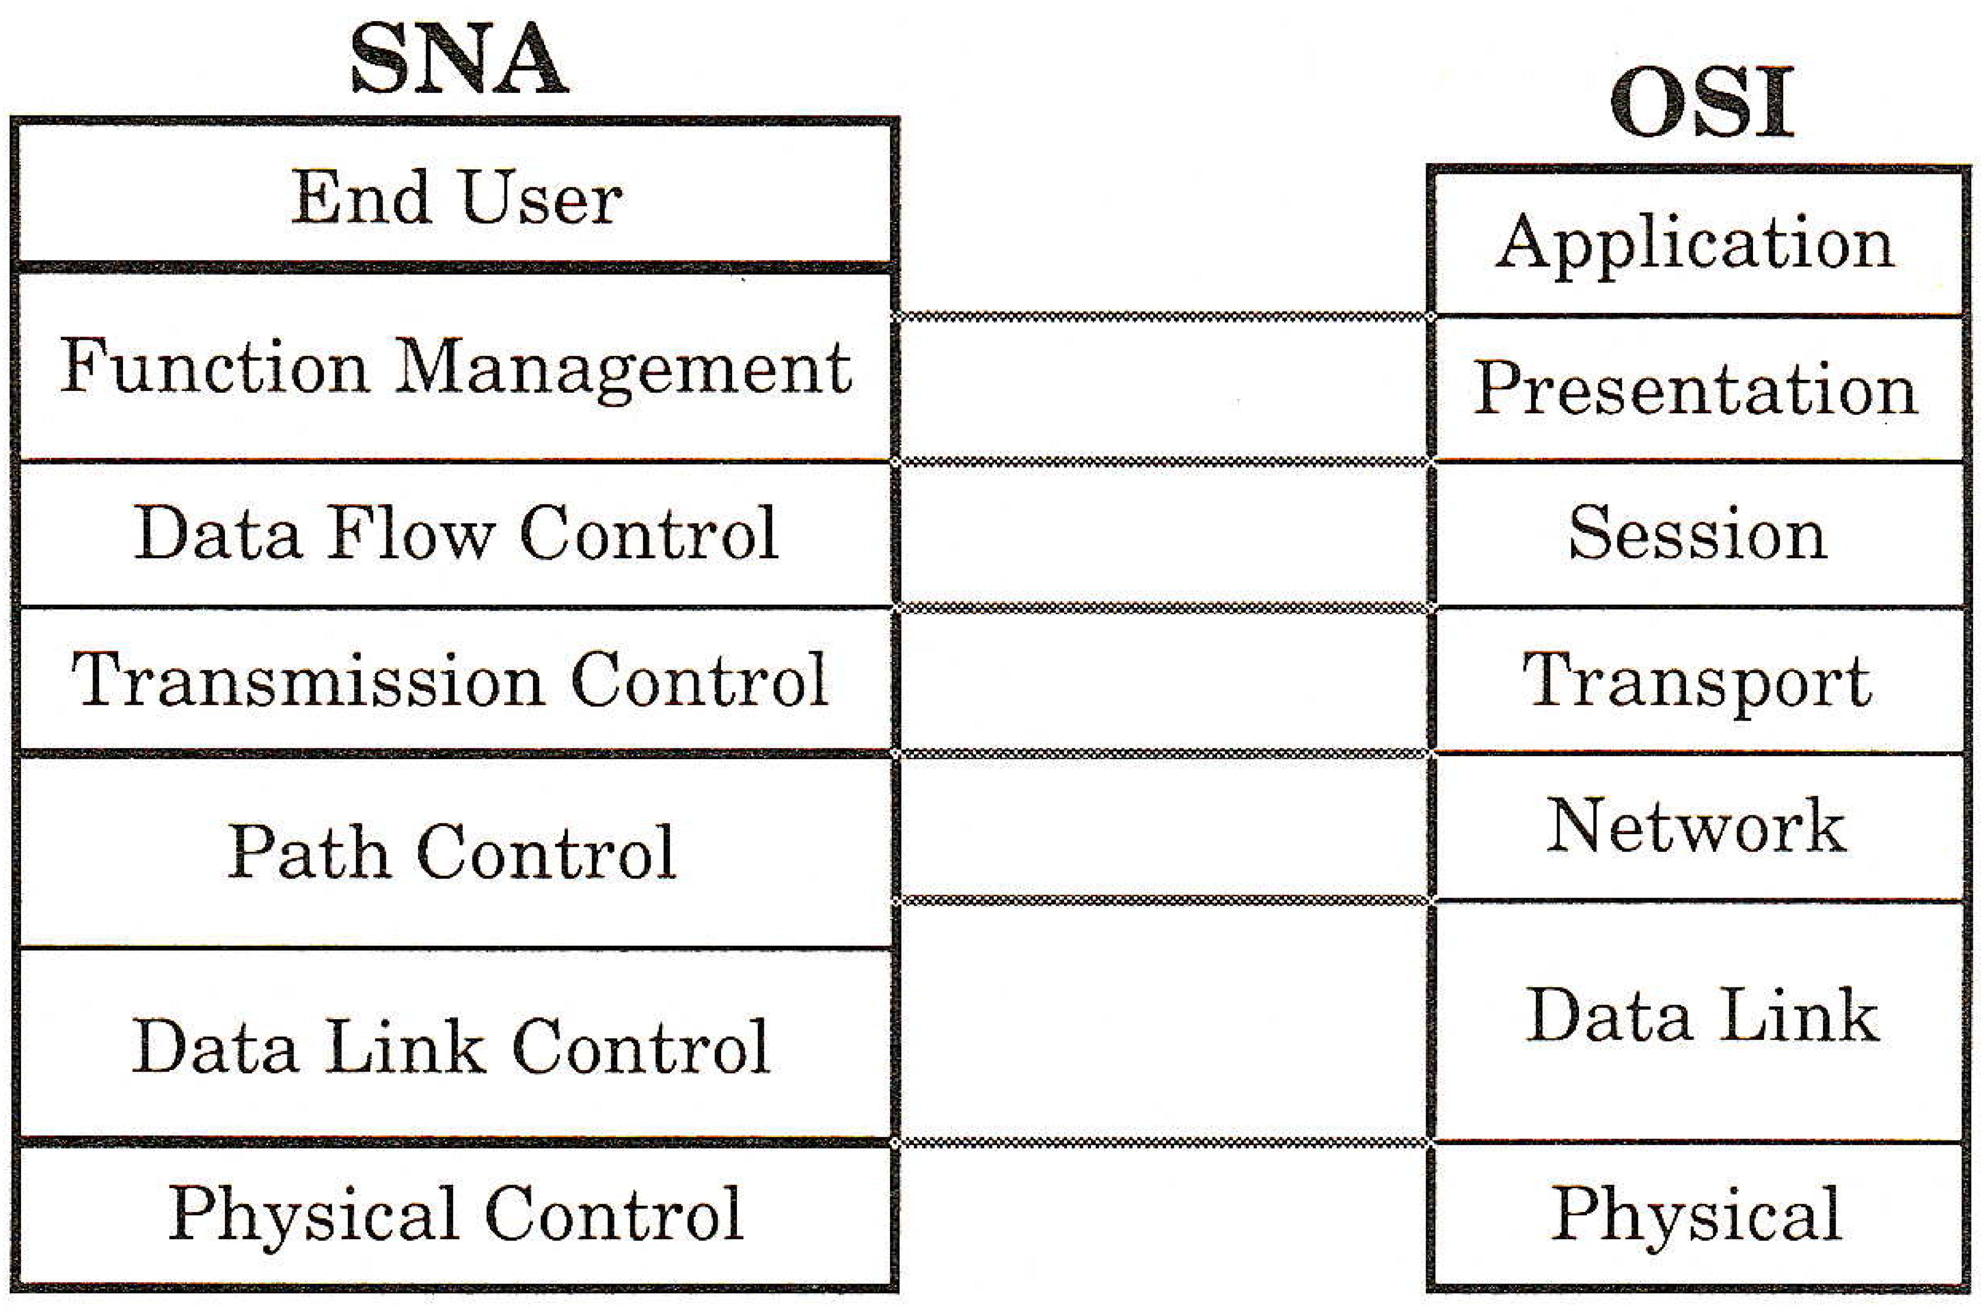

Figure 3: Layers of SNA and OSI

Early SNA internetworking

In the beginning, SNA protocols were transferred exclusively over serial lines using the *Synchronous Data Link Control* (SDLC) protocol. Saying that you used SNA assuredly also meant that you used SDLC—the two terms were often uttered in the same breath—SNA/SDLC.

The higher layer SNA protocols always assume a reliable, connection-oriented service at the data link layer. This guaranteed delivery mechanism at the lower layers greatly simplifies the upper layer SNA protocols. Unlike TCP, there is no need for the upper layers of SNA to maintain retransmission timers or compute checksums. If there are difficulties encountered with transmitting data between two SNA nodes, the lower layers detect the problem and notify the upper layers.

In traditional SNA, SDLC provides this reliable circuit for data exchange. Early in the life of SNA (circa 1980), IBM found the need to transport SNA protocols across existing X.25 [12] networks, especially in Europe, Canada, and Japan. On the surface, the X.25 *virtual circuit* provides the appearance of a single data link with end-to-end control much like SDLC. However, there was not a perfect match between the services required for the bottom layer of SNA (i.e., Path Control) and the services provided by X.25. While the HDLC information frames within X.25 could be used to carry the contents of the SDLC information frames, SDLC control frames presented a bit of a problem.

QLLC

The solution was a special purpose component inserted between Path Control and layer 3 of X.25 to provide this SDLC-like end-to-end control. The SDLC control frames were carried across the X.25 virtual circuit inside X.25 packets with the data qualifier bit turned on. As a result, this special component for transporting SNA across X.25 was called *Qualified Logical Link Control* (QLLC).

QLLC represented the first form of protocol wrapping for SNA. The architecture of two SNA nodes communicating via X.25 virtual circuits using QLLC is shown in Figure 4 on the next page.

continued on next page

SNA Internetworking *(continued)*

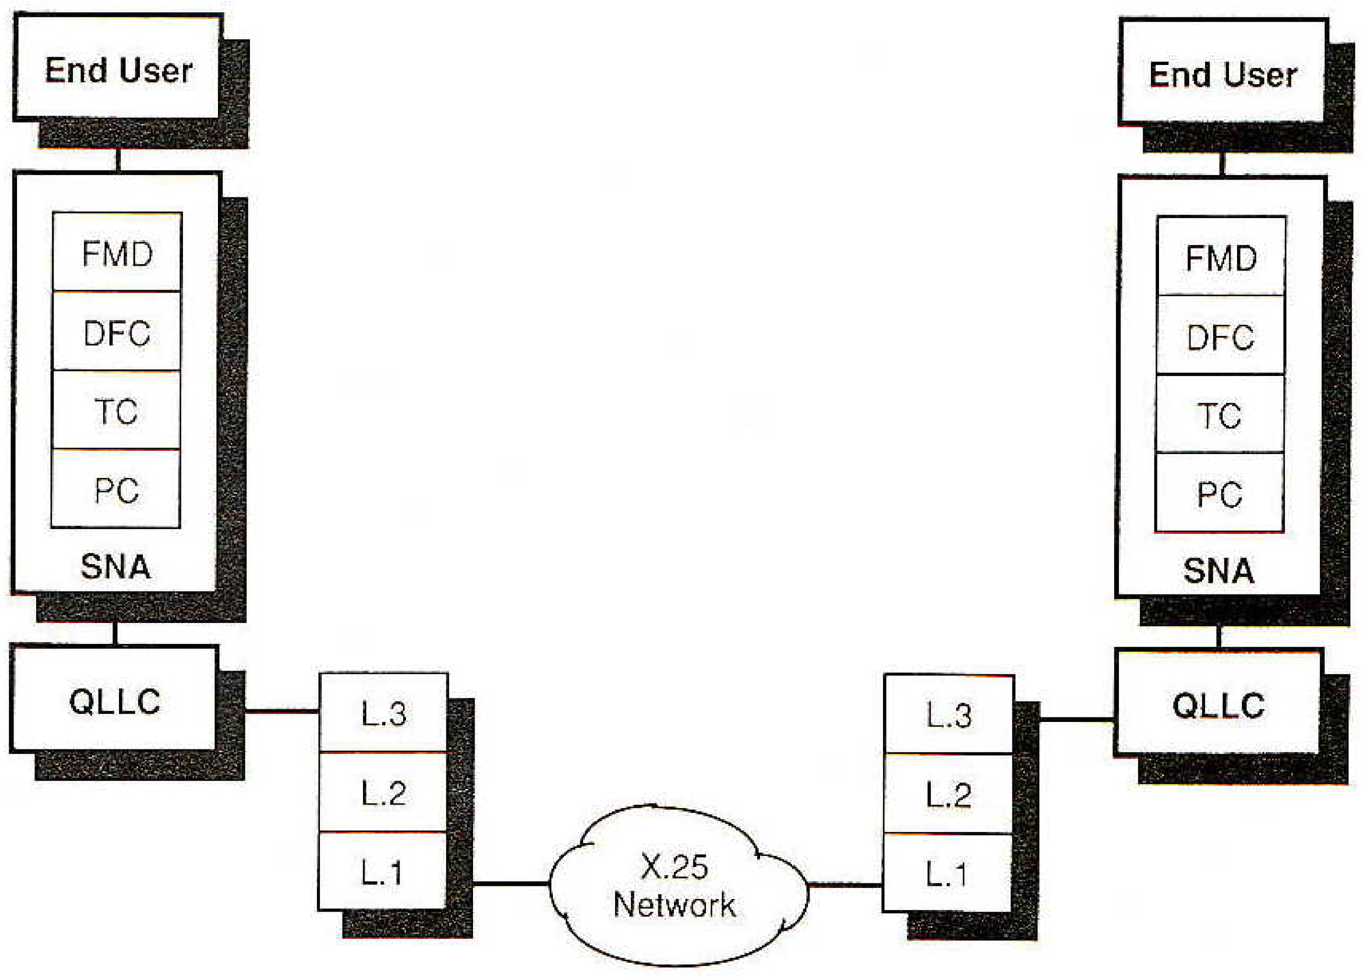

Figure 4: SNA Over X.25

Today's multiprotocol routers

All of the SNA interoperability solutions available in today's multiprotocol router market can be placed into one of the following two categories:

- *Data Link Control services:* This rather vague term includes not only traditional layer 2 bridging (such as source route bridging) but also the transformation of one data link protocol into another.
- *SDLC tunnelling:* SDLC frames are encapsulated inside another protocol (usually TCP/IP or HDLC) for transfer of SNA traffic over a multiprotocol network.

Data Link Control services

This category of internetworking services is the accommodation of LAN internetworking beneath SNA. These services operate exclusively at the data link control layer and cover bridging as well as data link transformations. Bridge technology that is appropriate for SNA-based LAN internetworks are *source-route bridging* (SRB), *remote source-route bridging* (RSRB), *source-route transparent* (SRT) and *translational bridging.* A network illustrating all of these bridging components is shown in Figure 5.

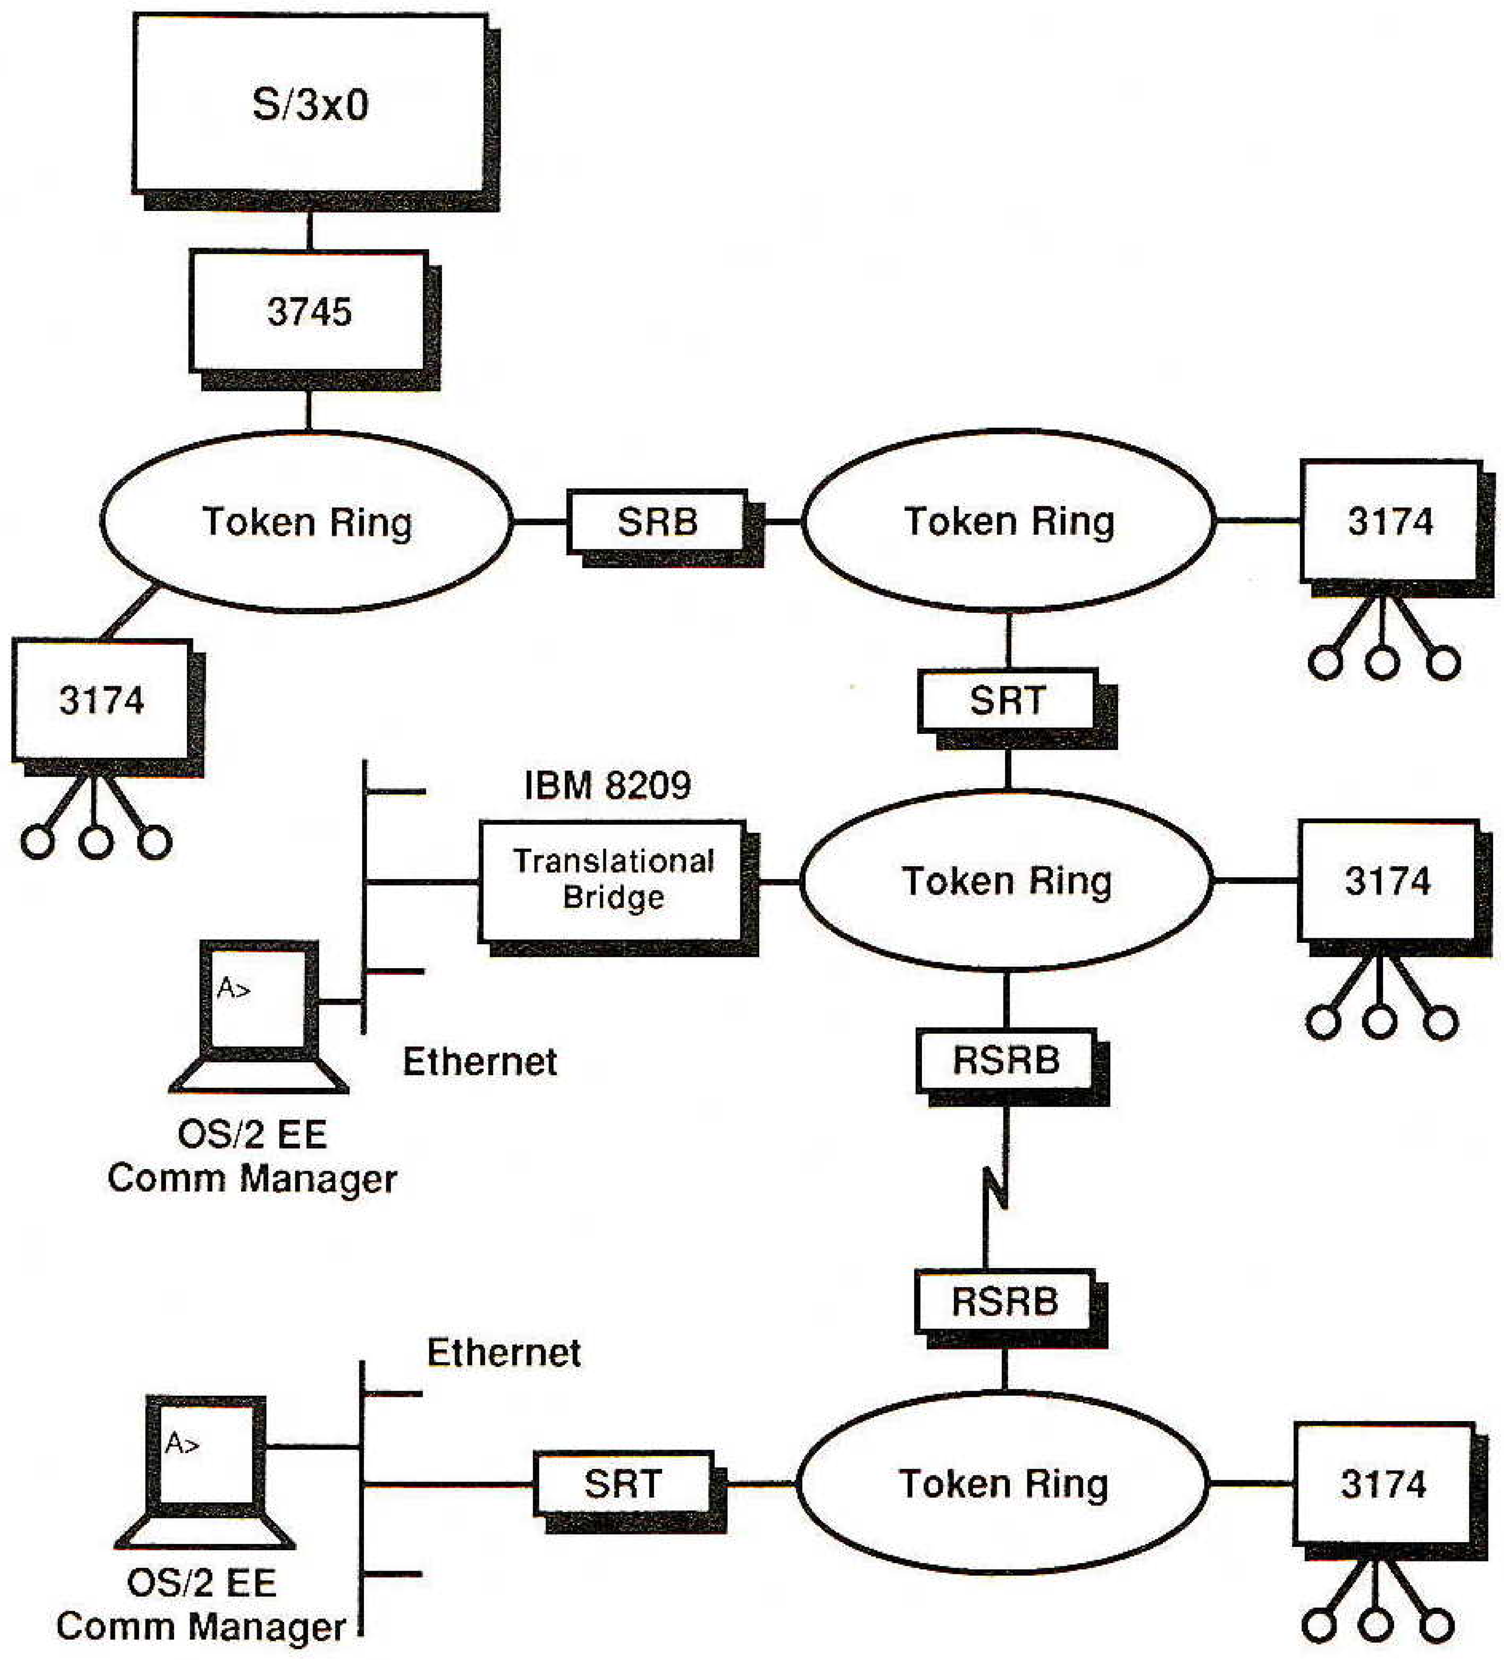

Figure 5: SNA In a Bridged Environment

While IBM was promoting its proprietary source route bridging algorithm as the preferred method of interconnecting token ring LANs, most other vendors were transparently bridging Ethernets using IEEE 802.1d's spanning tree algorithm. A good discussion of the pros and cons of source route bridging versus transparent bridging can be found in Tanenbaum [5]. Regardless of whether transparent bridging is better or source-route bridging is better, it is a fact of life that today's bridged token ring networks are source routed.

A variety of vendors including IBM, CrossComm, Wellfleet, Cisco, and Proteon provide not only source route bridges but also the capability of extending the source routing across a wide area network using appropriate wide area transport services such as 56 Kbps or T1. However, extending source routing across a wide area network has some well-known limitations.

LLC2

As indicated earlier, SNA requires a reliable, connection-oriented data link protocol. The data link protocol used on token ring LANs is therefore the connection-oriented variant of IEEE 802.2's *Logical Link Control* known as *LLC Type 2* or LLC2. LLC2 was originally designed to be a LAN protocol with no intentions of it being extended across a wide area network. Some of the timers associated with LLC2 expire when its frames incur the delay of a wide area network. The problem is exacerbated when LLC2 frames are encapsulated inside packets of a wide area network protocol that uses nondeterministic routing algorithms such as those found in TCP/IP. This particular problem is being solved by Cisco through local termination of the LLC2 sessions at the router. In this way, LLC2 sessions remain on the LAN and are not subject to wide area network delays.

Premature timer expiration is only one problem encountered with source routing. Others are Route Discovery storms and hop count limitations.

There are a couple of popular variants of translational bridging: Ethernet-to-token ring and SDLC-to-LLC2. A good example of Ethernet-to-token ring is IBM's 8209 bridge shown in Figure 5.

In recent months, several vendors have announced—and in some cases have already released—products that support a rather unique form of data link transformation: conversion of SDLC to LLC2. Unlike SDLC tunnelling which uses SDLC on both sides of the internetwork, with SDLC-to-LLC2 conversion, you see SDLC on one side of the *brouter* (concatenation of "bridge" and "router") and LLC2 on the other. This conversion makes a lot of sense to many IBM customers who have older SDLC-connected gear that cannot be upgraded to support token ring. By converting SDLC to LLC2, these older devices can gain ready access to token ring-based systems such as the newer IBM front end processors.

Technically, this conversion process is very straightforward since both SDLC and LLC2 are connection-oriented data link protocols and they share a common ancestor: ISO 3309's HDLC. An example of this form of SNA internetworking is shown in Figure 6 on the next page.

SDLC Tunnelling

Most of the vendors of multiprotocol routers either currently ship or have announced SDLC tunnelling products. SDLC tunnelling (also called *synchronous passthrough*) is the encapsulation of SDLC frames inside the packets of another protocol (usually TCP/IP or HDLC) and the transport of those encapsulated frames across a multiprotocol backbone network. The obvious result is that both SNA and non-SNA protocols can share the same network.

continued on next page

SNA Internetworking *(continued)*

Multiprotocol vendors have chosen this approach because of the extreme complexity of routing SNA in its native form along with the uncertain direction that SNA will take over the next few years. Therein lies the elegance of SDLC tunnelling: it is completely insensitive to the form of SNA that it is encapsulating. Not only do changes to SNA not affect the tunnel but the routers that support this capability do not care what kind of SNA nodes are at either end of the tunnel. To illustrate this property, Figure 7 shows two SDLC tunnels: one between a communications processor (a PU 4 device) and a cluster controller (a PU 2 device) and the other between two communications processors. The SDLC tunnelling logic is the same in both cases.

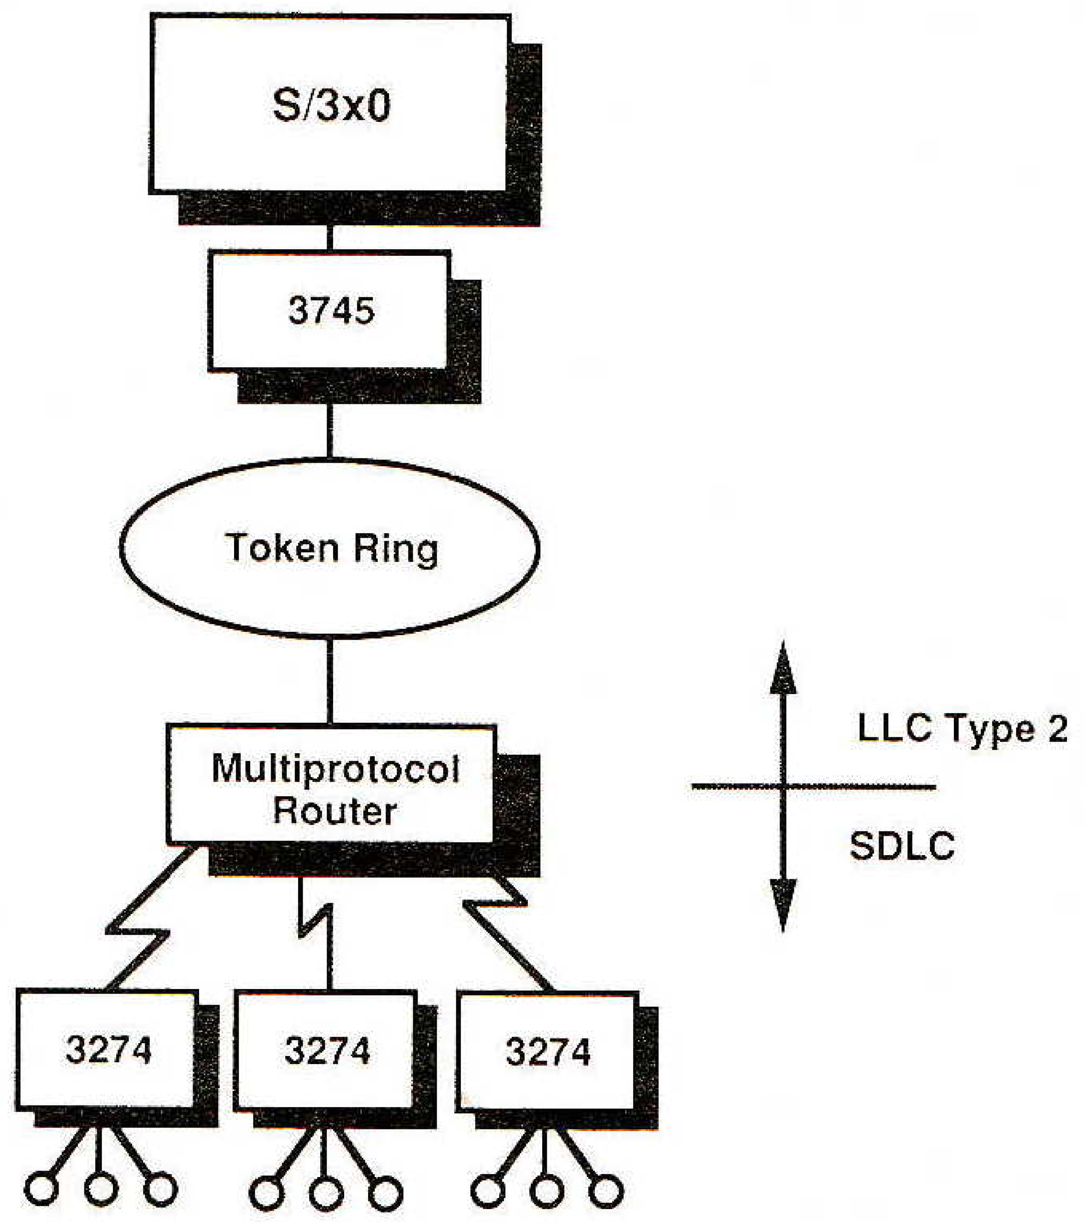

Figure 6: SDLC-to-LLC Type 2 Conversion

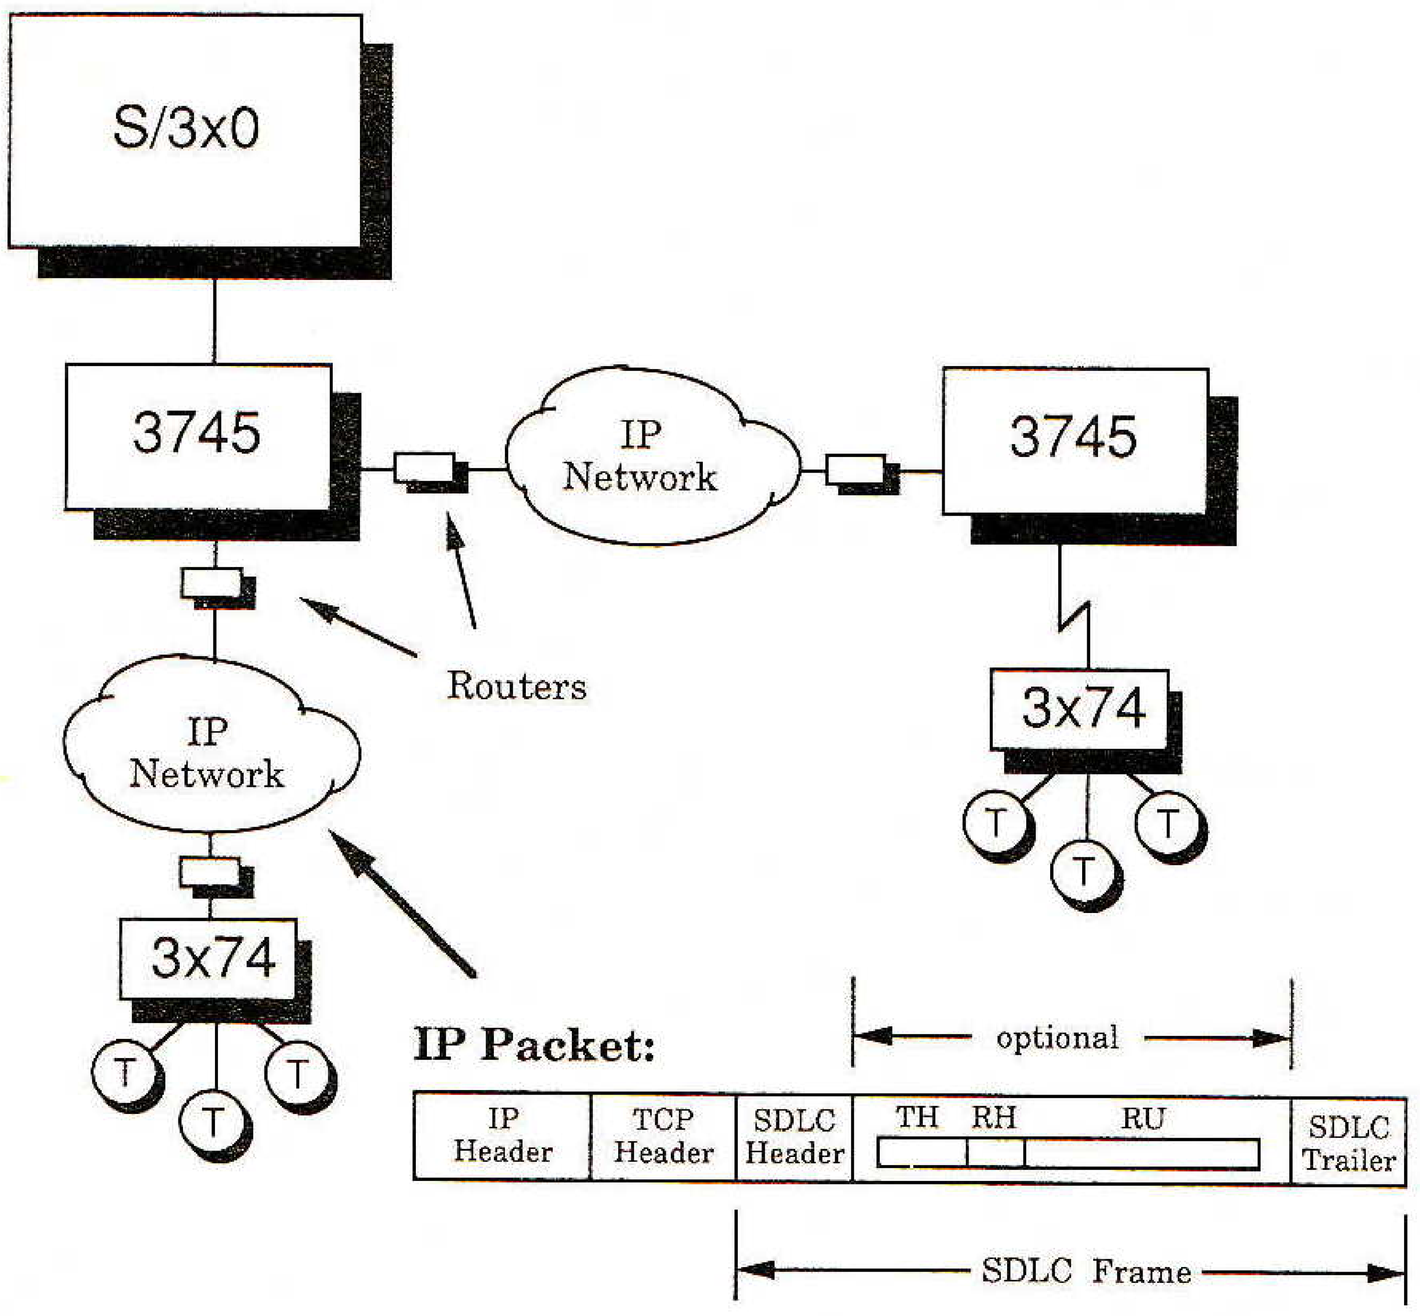

Figure 7: SDLC Tunnelling Across an IP Network

SDLC frames are encapsulated either inside another data link protocol such as HDLC (Cisco and Wellfleet) or LLC2 (CrossComm) or inside TCP/IP packets (Cisco and Proteon). Since SNA requires guaranteed end-to-end delivery of its data link layer and since it is the transport layer of the TCP/IP architecture that provides this level of service, SDLC is wrapped inside TCP packets rather than straight IP packets [2].

Serial tunnelling is not without its shortcomings. Below are three of the most significant problems with SDLC tunnelling:

- Lots of nonessential SDLC traffic in the form of "session keep-alives";
- End-to-end SDLC timers are sensitive to network delays;
- Variable response time.

Proxy Polling

The first was solved in Cisco's Serial Tunnelling (STUN) product through a feature known as *proxy polling*. Proxy polling (also known as *poll spoofing*) is a technique that eliminates most polls for data that is transmitted from the primary link station to the secondary. As long as the polls and the corresponding poll responses are non-productive (i.e., there is no data to be sent), the routers at each end of the serial tunnel can safely mimic the link station they are designed to replace.

Figure 8 demonstrates proxy polling. The router that is connected to the SNA node that implements primary SDLC performs the role of a secondary SDLC link station. Likewise, the router that is connected to the SNA node that implements the secondary link station acts as a primary SDLC.

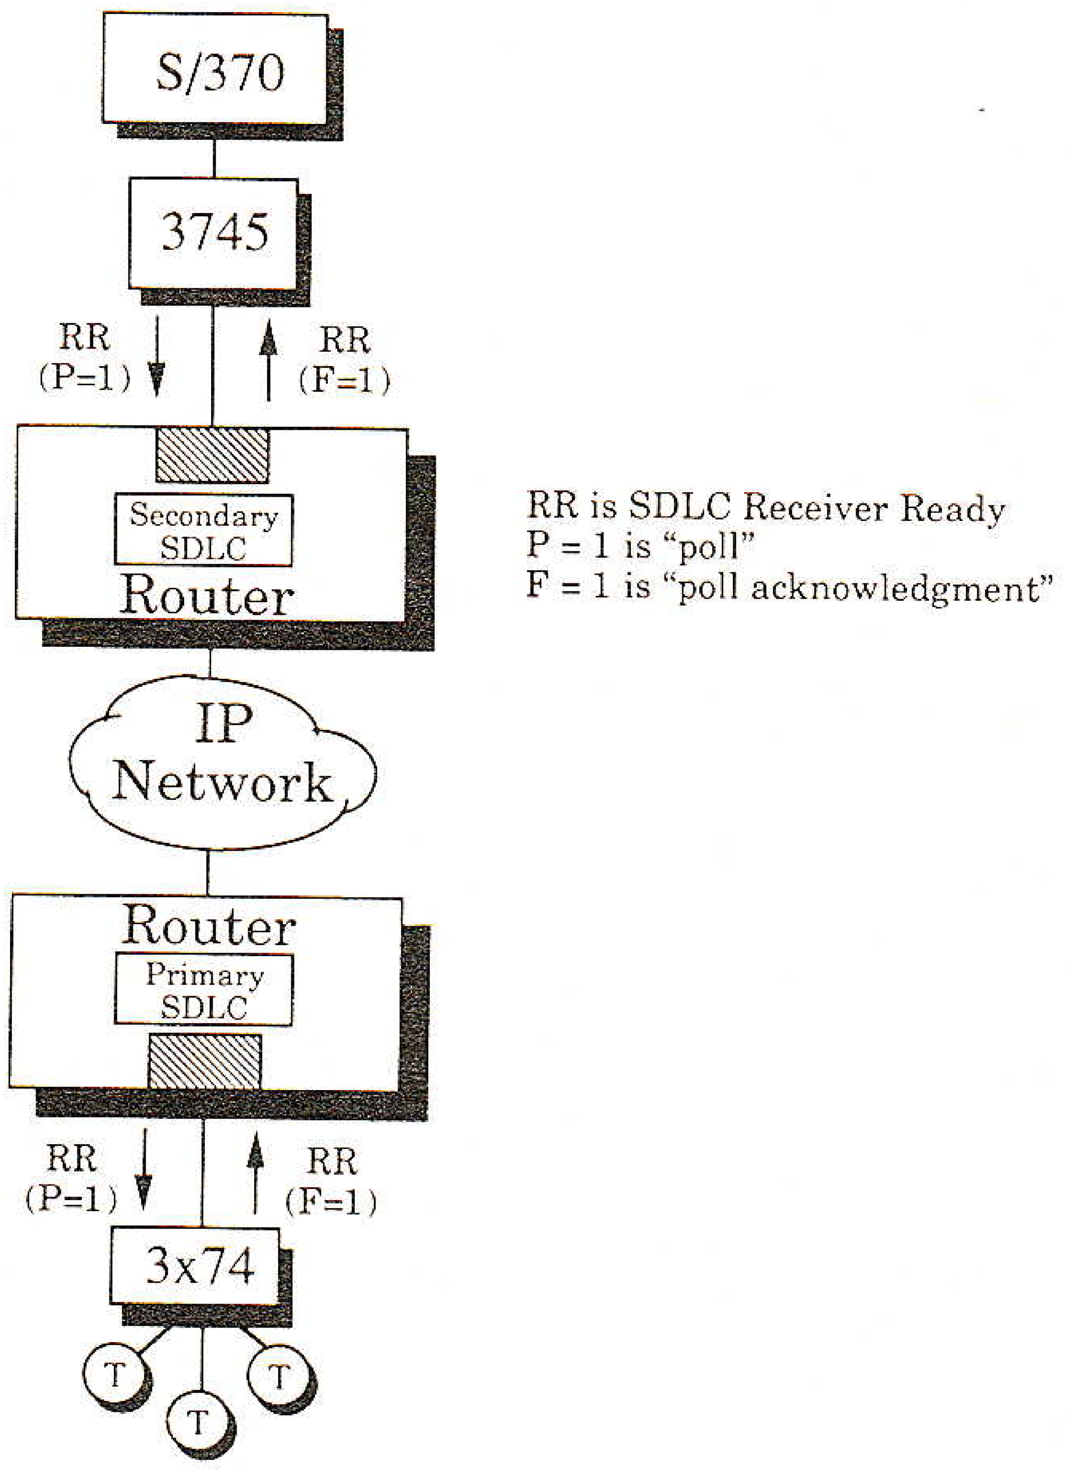

Figure 8: SDLC Tunnelling with Proxy Polling

While proxy polling reduces the amount of traffic on the intermediate network, it does not eliminate the problem with end-to-end timers. The problem that was described previously with LLC2 as it applied to remote SRB also raises its ugly head with SDLC tunnelling. The singular advantage that SDLC timers have over the LLC2 timers is that SDLC's timers are usually easier to increase since their definition is localized in the SNA network definition on the host (i.e., the sysgen).

There is a solution to the problem of SDLC timers and the solution is the same one as for remote SRB: locally terminate the SDLC session at the router. Even if the SDLC session is locally terminated and connections can now endure long delays, the situation with variable response time still exists. Tossing packets because of network congestion is unheard of in SNA, yet is standard procedure for IP networks.

continued on next page

SNA Internetworking *(continued)*

To an SNA network administrator, variable response time is worse than poor response time. Variable response time in an SNA network usually signals equipment failure so coping with the effects of dynamic routing and congestion control procedures will take some getting used to. Fortunately, the benefit of greater throughput of a multiprotocol backbone carrying encapsulated SNA traffic usually outweighs any concerns about variable response time.

SNA directions

To say that third-party vendors of SNA equipment have had a difficult time breaking into hierarchical SNA networks would be a gross understatement. The only niche that most SNA vendors have been able to exploit is in peripheral node emulation. In the early to mid 1980s, there were a variety of third-party manufacturers of IBM compatible cluster controllers. With the advent of local area networks, this niche has evolved in recent years into LAN gateways.

The only companies that have been even moderately successful at penetrating the SNA subarea network are NCR Comten and Amdahl [6]. This condition is mainly the result of IBM's stronghold on SNA subarea protocols. The closed, proprietary nature of SNA's Physical Unit Type 4 has resulted in today's distinct separation between SNA networking and what has become multiprotocol networking.

The last two years have not been kind to SNA. The chasm between open network architectures (TCP/IP and OSI) and SNA has resulted in vastly divergent capabilities and even greater differences in price/performance. With most vendors hopping onto the "open" bandwagon, IBM has been left alone to champion the cause of SNA and to play catch up with recent technical advances (e.g., Frame Relay [13] and FDDI [14]).

IBM freely admits that the statically-defined, hierarchical SNA typified in today's subarea networks is an anachronism and needs to be replaced [7]. Of course, they are advocating their own APPN architecture as the heir apparent. But this is not 1974 and IBM is not the only game in town any more. The world has solved the problem of easily connecting autonomous computing systems and APPN, while a great leap forward from its hierarchical predecessor, doesn't provide much real benefit to the networking world [8].

The reservations being expressed by networking pundits has not slowed down the rollout of APPN implementations. If anything, it has increased the tempo. Even though APPN has been available as a product since 1985, it was only implemented on the company's midrange computers: the System/36 and its successor, the AS/400. That condition changed in 1991 with the rollout of APPN on the 3174 cluster controller and OS/2 Extended Edition. IBM has indicated that VTAM will acquire APPN functionality in early 1993 [15] and the RS/6000-based multiprotocol router will have APPN capability either late in 1992 or early in 1993.

Multiprotocol SNA routing

The metamorphosis of SNA presents a significant challenge to vendors of multiprotocol routers: even though the current installed base is hierarchical SNA, IBM is moving swiftly to deploy APPN into all reaches of SNA-dom [10]. If a router were to implement PU 4 protocols and interoperate with IBM's 3745s, there might not even be a PU 4 to talk to in the market by the time it's completed :–). Even IBM has publicly stated that it does not plan to include PU 4 on its own multiprotocol router but does plan APPN for the second release in the middle of 1992 [9].

While IBM has done a good job of telling the world about the wonders of APPN, few IBM customers have yet to embrace it. In order for APPN to be successful, it has to be readily available from not only IBM but also third-party vendors. Four companies were supposedly working on APPN End Node implementations by virtue of getting the early specifications of the End Node protocol from IBM: Systems Strategies Inc., Novell, Apple Computer, and Siemens. None of them have announced an APPN product. This means that all APPN products today come from only one vendor: IBM. The lack of success that these third-party vendors have had with the simpler APPN End Node spells doom for APPN Network Node (the APPN node that does routing) even if the much-anticipated specifications were made available.

Never underestimate IBM's clout at marketing its products. But IBM's only hope of making APPN anywhere as successful as SNA has become lies in having third-party offerings of APPN available. IBM has said it will license APPN Network Node source code to other communications vendors sometime during 1992 [16]. If IBM hopes to stem the tide of customers abandoning SNA for non-SNA solutions, IBM must not only make APPN available on all of its own platforms but it must also provide a robust, portable implementation of APPN to the rest of the industry.

References

This book ——> *is reviewed on page 29.*

[1] Study performed by Computer Intelligence Corp., 1990.

[2] Clark, W., "Serial Tunnelling: Spoofing Your SNA Network," *SNA Perspective,* Volume 12, Number 10, October, 1991.

[3] *Communications for Cooperating Systems: OSI, SNA, and TCP/IP,* R. J. Cypser, Addison-Wesley Publishing Co., 1991.

[4] *SNA: IBM's Networking Solution,* James Martin, Prentice-Hall.

[5] *Computer Networks,* Second edition, Andrew S. Tanenbaum, Prentice-Hall Inc., 1988, ISBN 0-13-162959-X

[6] Kucharski, A., "IBM's Marathon Runner: The Communication Controller," *SNA Perspective,* Volume 12, No. 10, October, 1991.

[7] Clark, W., "The SNA Subarea Under Siege," *SNA Perspective,* Volume 12, Number 9, September 1991.

[8] Mohen, J., "APPN: The Networking Cornerstone that Isn't," *Data Communications,* Volume 20, Number 4, April 1991.

[9] "IBM Exec Shares Plans for APPN Router," *Network World,* Volume 8, Number 50, December 16, 1991.

[10] Guruge, A., "IBM Radically Reformulates SNA," *Data Communications,* Volume 20, Number 4, April 1991.

[11] Vair, D., "Components of OSI: X.25—the Network, Data Link, and Physical Layers of the OSI Reference Model," *ConneXions,* Volume 4, No. 12, December 1990.

[12] Malamud, C., "DECnet/OSI Phase V: Real OSI or Only Selected Interfaces?," *ConneXions,* Volume 4, No. 10, October 1990.

[13] Kozel, E., "The Cisco/DEC/NTI/StrataCom Frame Relay Specification," *ConneXions,* Volume 5, No. 3, March 1991.

[14] Wolter, M., "Fiber Distributed Data Interface (FDDI)—A Tutorial," *ConneXions,* Volume 4, No. 10, October 1990.

[15] "IBM Lays Out APPN Plan for Mainframes," *Network World,* Volume 9, Number 4, February 3, 1992.

[16] "IBM Launches Line of SNA-Ready Routers," *Network World,* Volume 9, Number 3, January 27, 1992.

WAYNE CLARK is the SNA Architect in the Engineering Department at Cisco Systems Inc. in Menlo Park, California. He has been working in the computer industry for 20 years and has worked on SNA implementations at Memorex, Ungermann-Bass, Communications Solutions Inc. (later part of 3Com), and Novell. His specialty is the implementation of LU 6.2 and NT 2.1 protocols. He was also the instructor of Interop's very first SNA tutorial in 1991. He has an M.S. in Computer Engineering from Santa Clara University. He can be reached on the Internet as: `wclark@cisco.com`.

Lessons Learned for OSI at INTEROP 91 Fall

by Susan Hares, Merit/NSFNET

Introduction

Thirty networking technology vendors worked together to provide a demonstration of OSI applications and network protocols over the Internet's Infrastructure for INTEROP 91 Fall. The National Agency networks participating in this event were NSFNET, ESNET, and the NASA Science Internet. The demonstration linked workstations on the convention show floor with systems in Europe, the United States and Australia (see Figure 1).

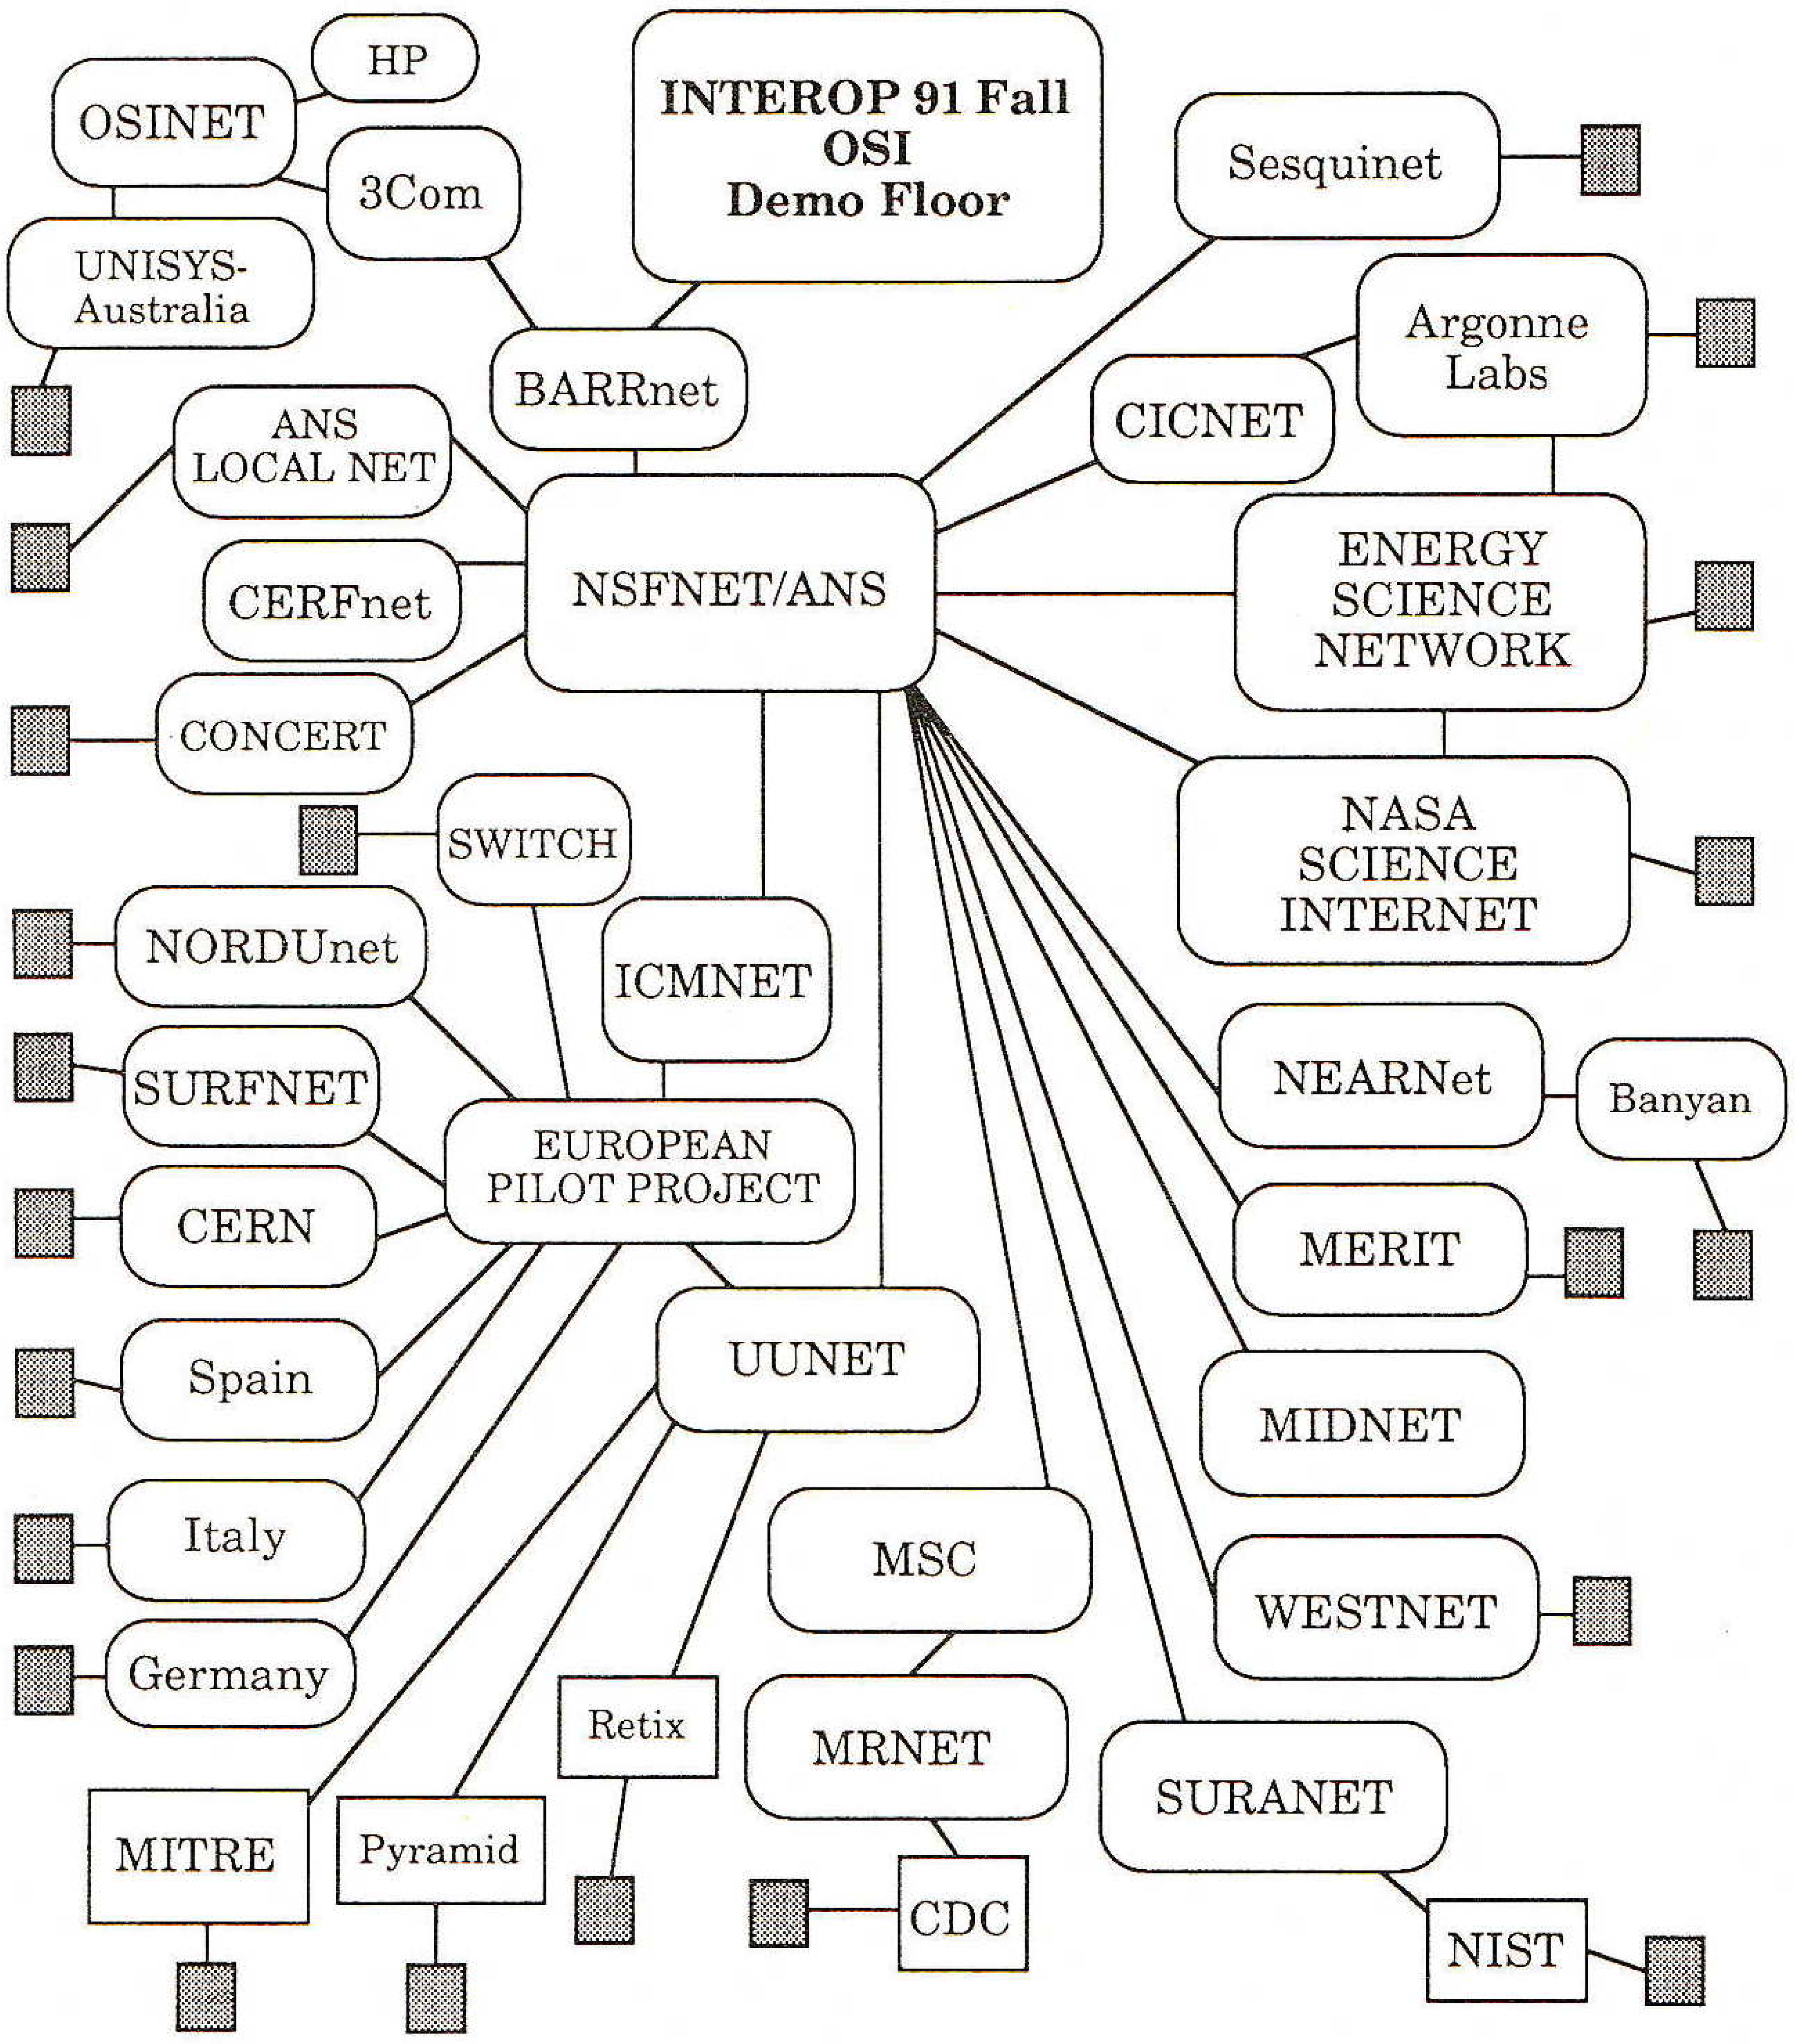

Figure 1: OSI Infrastructure in the Internet

This INTEROP 91 Fall demonstration showed how cooperation from many network service providers and vendors can make OSI applications over the Internet a reality. The lessons learned from this demo pave the way for continuing OSI traffic in the Internet. OSI vendors who participated want to continue to use the Internet during the next year.

Obstacles

The most painful lessons at INTEROP 91 Fall were the obstacles to showing the OSI applications. The OSI Demonstration booth encountered three types of problems:

- Problems setting up the physical network,
- Problems setting up and debugging a multi-protocol network, and
- Problems configuring OSI routers and debugging OSI related problems.

Setting up a huge multi-protocol network in only a few days is close to impossible. Only extensive planning, testing, and excellent cooperation from the Internet community (vendors and network service providers) has allowed Interop, Inc. to accomplish this difficult task year after year. Every booth encountered some problems due to the physical network set-up and multi-protocol network debugging. These problems are normal when setting up this kind of tradeshow. In addition, the OSI Demonstration booth encountered difficulties due to:

- Problems with the set-up of OSI routing of *Connectionless Network Layer Protocol* (CLNP) packets via static configurations, and
- Lack of OSI network debugging tools on every machine.

Sometimes OSI application traffic flowed from the show floor to Europe, but not between booths on the show floor!

What happened at INTEROP 91 Fall

OSI demonstrations on the INTEROP 91 Fall show floor included OSI Vendor booths and the collaborative OSI Demonstration booth (see Figure 3). ISO's *End System to Intermediate System* (ES–IS) protocol was used between multiple hosts (end systems) and routers (intermediate systems) from many different vendors. ISO's *Intermediate System to Intermediate System* (IS–IS) protocol was shown in the OSI Demonstration booth. IS–IS is an ISO intra-domain routing protocol which is similar to *Interior Gateway Protocols* (IGPs) in the TCP/IP suite. NSFNET has used an implementation of the IS–IS routing protocol adapted for IP since 1988.

Applications

The applications showcased at the OSI demonstration booth and vendor booths included four major OSI applications: *X.400* (Messaging), *File Transfer Access and Management* (FTAM), *Virtual Terminal* (VT), and *X.500* (Directory Service). X.500 was demonstrated over both IP and CLNP showing that OSI applications do not have to be limited to OSI lower layer stacks. Figure 2 shows Internet applications which have some of the functions of these four major OSI applications.

OSI application:	*Internet application(s) with some of the same functions:*
VT	Telnet
FTAM	FTP
X.400	SMTP
X.500	none *

**Note: The TCP/IP protocol suite has no protocol that provides the distributed directory service that X.500 provides. X.500 is being used over TCP/IP in the Internet.*

Figure 2: OSI applications compared to TCP/IP applications

IDRP

A prototype of the ISO *Inter-Domain Routing Protocol* (IDRP) was demonstrated. IDRP provides a means of passing OSI routing information between domains and applying policy filters to that routing information. An IP protocol which provides some of this functionality for IP is the *Border Gateway Protocol* (BGP). IDRP is in the second stage of the development process as an ISO standard. The development of the IDRP prototype provided a great deal of feedback on this ISO standard (CD 10747) to the US committee working on it. The IDRP prototype passed traffic between nodes in the OSI demonstration booth, the IBM booth, and nodes on the ANS/NSFNET T3 test network. IDRP was developed by Dave Katz of Merit, and further details can be obtained from Merit.

continued on next page

Lessons Learned for OSI at INTEROP *(continued)*

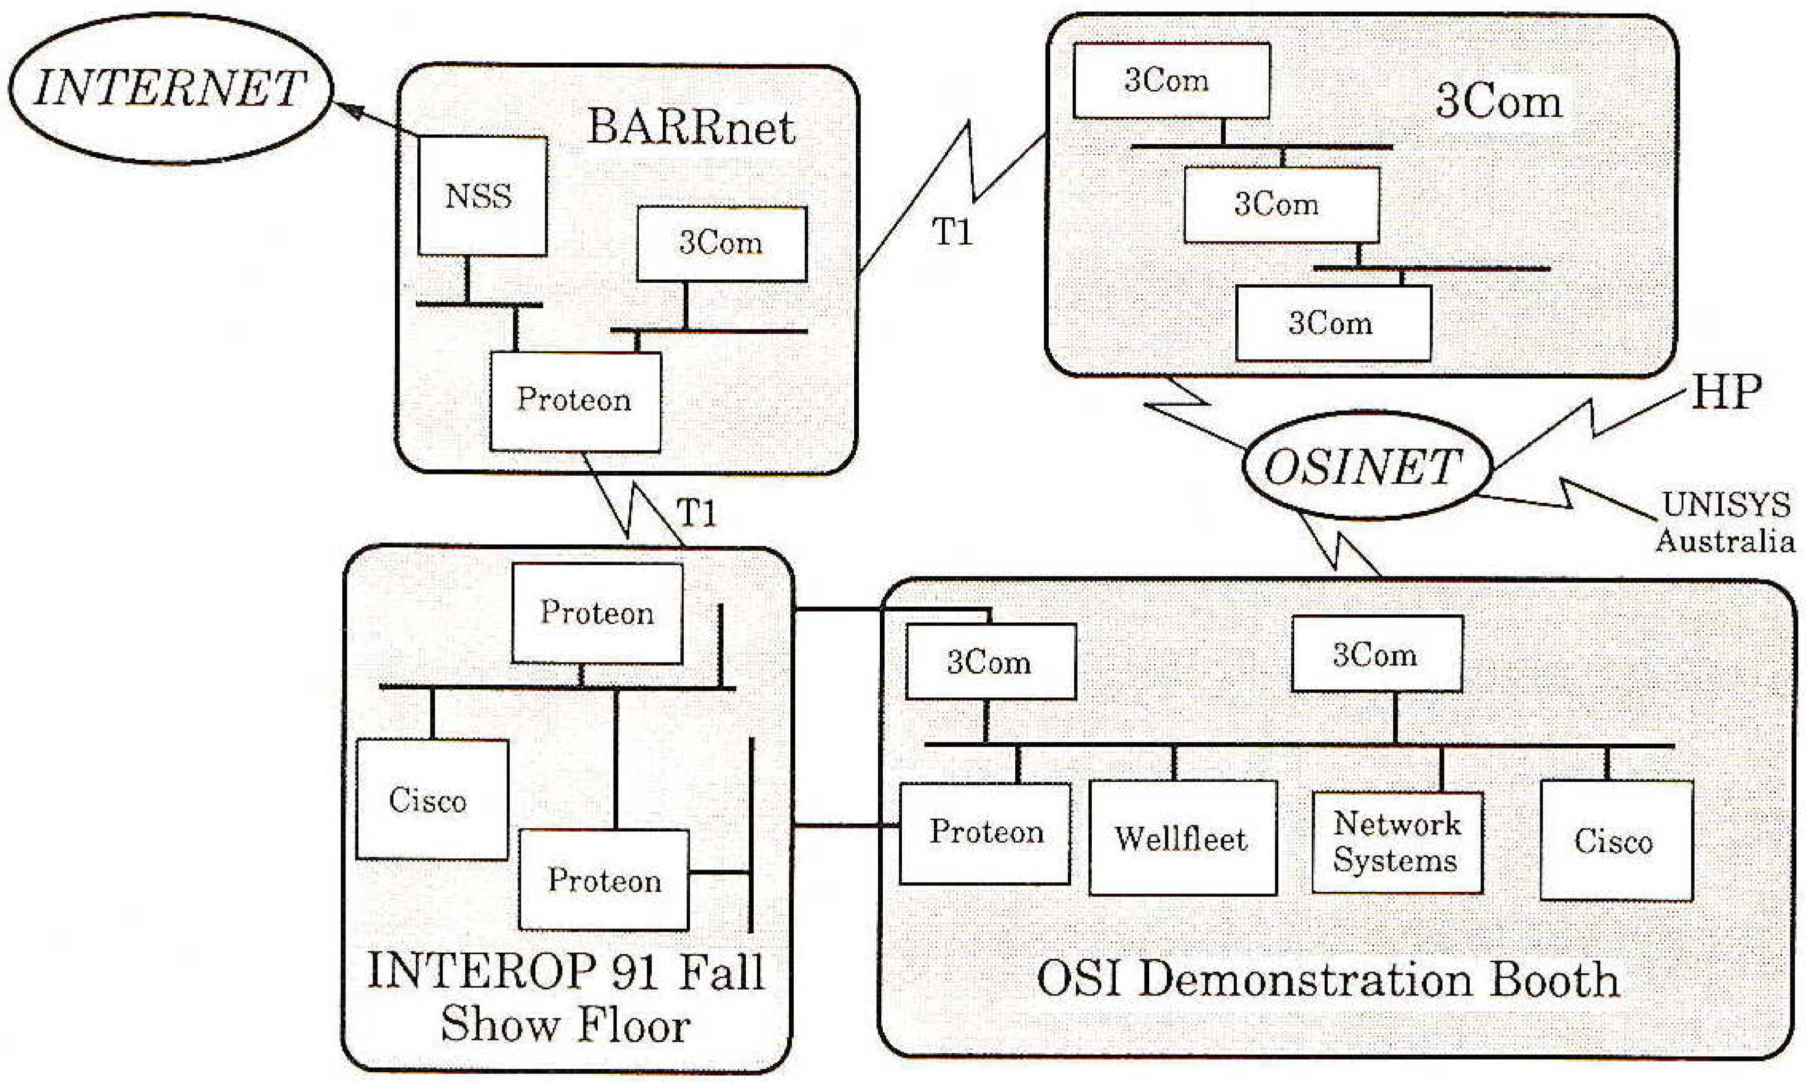

Figure 3: OSI at INTEROP 91 Fall

Lessons

• *Lesson 1: People make OSI Work.* Many talented people made the OSI Infrastructure demonstration happen. I am convinced that the Internet networks work and advance because talented individuals push and push until the technology advances. Companies support technology advances by putting the punch behind their people.

Cyndi Jung (3Com) helped me organize the OSI demonstration booth and Internet testing. She spent countless hours working on the OSI Hot Stage and the Backbone Hot stage. The router vendors (Cisco Systems, Proteon, 3Com, Wellfleet, Network Systems, DEC) spent extra time helping out the network, Hot Stages and IS–IS testing. Two especially hard workers were Paulina Knibbe (Cisco) and Ed Stern (Proteon).

End system vendors worked with Network service people to set-up the OSI application demonstrations. Night after night Charlie Alberts and John Davis (Banyan), Kevin Jordan and others from CDC, Eva Kuiper (HP), and other vendors tested FTAM, X.400 and X.500 across the Internet. Cathy Wittbrodt, Arlene Getchell and Tony Genovese of ESNET made the ESNET networks and OSI applications work. Juha Heinänen and lots of people from the RARE-WG4 CLNS project helped connect Europe to this demonstration. Linda Winkler (Argonne Labs), Alan Clegg (CONCERT), Mark Knopper (Merit), Walt Lazear and John McGuthry (MITRE), Doug Montgomery (NIST), Cathy Fouston and Bill Manning (Sesquinet) got more OSI applications working across the Internet. The roll call for the network path includes many of the people who make IP a reality today: Vince Fuller and Ron Roberts (BARRnet), Mark Oros (ICMNet), Tim Salo and Jeff Wabik (Minnesota Super Computer Network), John Curran (NEARnet), Dave O'Leary (SURANet) and Andrew Partan (UUNET).

Some of the networks companies that lent their people, equipment and push to the OSI Infrastructure demonstration included: Alcatel TITN, 3Com, ANS, Argonne National Laboratory, AT&T, Banyan, BARRnet, CERFNET, CERN, CICNET, Cisco, Control Data Corporation, CONCERT, Digital Equipment Corporation, ESNET, Frontier Technology, HP, IBM, ICMNet, INFN (Italy), networks in Spain, networks in Germany, Merit, MIDNET, MITRE, Minnesota Supercomputer Network, MRnet, NASA Science Internet, NORDUnet, NEARnet, Network Systems, Novell, NIST, OSINET, Pyramid, Retix, SURANet, SWITCH (Switzerland), Tandem, UNISYS, UNISYS-Australia, The Wollongong Group, and Westnet.

One thing that helped harness these people were numerous conference calls provided by MCI.

• *Lesson 2: Build on the Past.* The OSI Infrastructure demonstration is the culmination of years of work. The idea for the Infrastructure demonstration was conceived in mid June of 1991 as a milestone for the long term work in the US Internet. The extension of OSI application traffic to some 30 networking technology vendors over a good portion of the Internet took place within 3 months. The rapid deployment of this Infrastructure was due to:

- Past work in OSI by Pilot Projects in Europe and the US, and
- Outstanding work by each of the networks and companies participating in the demonstration.

OSI support in routers and end systems has matured a lot in the past year. Pilot Projects in European and the US have caused some of these products to mature. The Pilot Projects have tested products, reported bugs, and suggested improvements to user interfaces and product capabilities.

Pilot Project history

NSFNET demonstrated a prototype implementation of CLNP at INTEROP 89 in September of 1989. The T1 NSFNET has been capable of routing CLNP since August of 1990. During August of 1990, Merit exchanged CLNP packets with sites in Europe as part of a project involving European Pilot Project sites and the NSFNET. During INTEROP 90, CLNP packets were exchanged between systems on the show floor and systems in Europe which were a part of the European RARE-WG4 CLNS Pilot Project.

During the time period between October 1990 and April 1991, MITRE and other US companies exchanged information using OSI applications (FTAM and X.500) with systems in Europe participating in this European RARE-WG4 CLNS Pilot Project. In April of 1991, two US sites—Merit and MITRE—successfully sent files between US systems on the Internet over a pure OSI stack using CLNP for the network layer protocol. The network pathway between these hosts was set up with the help of MichNet, UUNET, and NSFNET.

The Energy Science Network (ESNET) has been working toward providing OSI within its backbone since early 1991. During June of 1991, ESNET was capable of routing CLNP packets to several sites. By September, ESNET could route CLNP packets in all routers on ESNET's backbone. By October 1991, several ESNET sites had hosts that exchanged information using OSI applications (FTAM, X.400, X.500) over a pure OSI stack using this CLNP pathway.

• *Lesson 3: Test Everything You Can.* Testing for the OSI demonstration booth was continuous from mid-June (when the idea was conceived) until INTEROP 91 Fall. While two "Hot Stages" were involved for this demo, most of the network testing for the OSI infrastructure demonstration took place outside of the Hot Stage time periods. The following testing was done:

- Pre-Hot Stage Internet Set-up and Tests
- NIST IS–IS Test Lab
- OSI Hot Stage
- Backbone Hot Stage
- Internet Testing of Applications

continued on next page

Lessons Learned for OSI at INTEROP *(continued)*

Need for a testbed

All of this testing took time, but pointed out the need for an on-going testbed. The IETF (Internet Engineering Task Force) *Network layer OSI OPerational* (NOOP) working group will be investigating how to organize a testbed for OSI. This testbed needs to have all routers running CLNP, and IS–IS within some routers. This national testbed needs to have end systems actively running OSI applications. Pieces of this testbed need to be linked together using the production Internet. OSI applications need to run between systems on the testbed as well as on systems in the production Internet.

Problems on operational networks can be reduced by putting new router software into testbed nodes, and passing OSI and IP application traffic over these routers. A continual testing of router and application software will lessen the time needed for INTEROP testing. An added benefit of a national testbed is that demonstrations of OSI applications over the Internet could be staged quickly.

Testing via the Internet

Setting up each Internet connection to a site running an OSI application involved the following steps:

- *Step 1:* Get the permission of one or more of network service providers to pass CLNP packets.
- *Step 2:* Set up routers in the networks to route CLNP packets.
- *Step 3:* Set up OSI applications on hosts.
- *Step 4:* Test OSI applications over these networks.

During June, July and August, a great number of the sites on the Internet followed these four steps and got OSI application traffic flowing across the Internet.

NIST testing of IS–IS

The *National Institute of Standards and Technology* (NIST) hosted a week of pre-INTEROP 91 Fall dynamic router interoperability testing, August 12–16. The open lab was part of NIST's *Cooperative Laboratory for OSI Routing Technology* program. Testing addressed the Draft International Standard specification of the IS–IS protocol (DIS 10589) and the operation of the IS–IS protocol in a live multi-vendor Intermediate and End System environment. OSI routing was tested over Ethernet. Vendors participating in the IS–IS testing included 3Com, Digital Equipment Corporation, Proteon and Wellfleet. This testing allowed the OSI Hot Staging to use the IS–IS protocol to support OSI applications. Breaking off routing protocol testing from OSI application testing greatly improved the OSI Hot Staging efforts.

Booth Hot Staging

In the last week of August 1991, a Hot Stage was held for the OSI demonstration booth. A T1 link from BARRnet to the OSI Hot Stage was essential for the success of the testing between vendors and OSI infrastructure demonstration. Full fledged Internet access allowed the vendors to exchange mail and obtain software changes from their company. Due to the full fledged Internet access not all team members working on an OSI product needed to attend the Hot Stage. Several experts from OSI vendors used the Internet to login to Hot Stage systems. Without leaving his/her desk at work, an expert could examine problems and try out solutions. Companies provided expertise without the cost of sending an additional person to the Hot Stage.

Application traffic for the OSI Infrastructure demonstration flowed across the T1 link to BARRnet. Due to this T1 link, Infrastructure demonstration throughout the Internet could be debugged from the Hot Stage area.

Backbone Hot Staging

INTEROP 91 Fall Backbone Hot staging started during the last week of August and continued throughout September. The backbone hot stage had the challenging task of putting IP, OSPF, CLNP on interfaces with FDDI, serial lines, Ethernet, and ISDN. Most networks only use a fraction of these protocols. The combination of these protocols in a multi-vendor environment broke new ground. The backbone Hot Staging showed that not only OSI, but other protocols can benefit from a national testbed for router software.

Another lesson from the Hot Staging is the need for good Internet connectivity to any demo activity. During the early stages of the Backbone Hot Staging, the Internet connection was not a full T1 link. Without this T1 link, access to Internet resources for exchange of mail or code updates was slow and problems took longer to fix. Experts from router vendors could not quickly access the backbone from their desk as they could with the OSI Hot Stage.

• *Lesson 4: OSI addressing needs full X.500 service.* Keeping up with the changing OSI addresses for each router and end system was quite a chore. The X.500 work for placing network addresses or host addresses in a global X.500 directory is not complete. While there is progress being made on these issues in several pilot projects, most of the addresses for the routers and end systems for INTEROP were kept in files. Managing these files took a great deal of my time, and needs to be improved for future Internet work.

An interim place to register network addresses and OSI application addresses needs to be in place while X.500 work continues. As a result of the INTEROP work, the mail group `osi-pilot@merit.edu` will help work with this addressing nightmare.

Formats

Router and end system vendors use different formats to express network addresses (*Network Service Access Points* [NSAPs]) and OSI application addresses. A common format for expressing these addresses would greatly speed debugging of network problems. Network debugging tools would only have to function on *one* format, not several.

• *Lesson 5: Routing Protocol (IS–IS) is much better than Static configurations.* The OSI Demonstration booth ran IS–IS between routers within the booth, BARRnet, and 3Com. The IS–IS routing protocol handled loss of links, and automatically switched to a backup route. The routers involved in the IS–IS protocol (3Com, Proteon, and Wellfleet) passed lots of traffic to the Internet.

The INTEROP 91 Fall show floor backbone routers ran static route configurations for CLNP. While the software on these routers ran well, setting up these configurations for the large show network was difficult. The lessons the Internet at large has learned about static routes for IP were re-learned as we tried to use static routes for OSI. The worst stress on static routes came when the INTEROP show floor team changed router vendors for some of the show floor routers (due to non-OSI protocol issues) on the day before the show opened. This last minute change required re-doing a lot of OSI static routing configurations in the new routers within hours of when the show opened.

Even with improved tracking of router configurations by the INTEROP NOC, the sheer amount of static configurations made life difficult for the show floor network. In fact, some problems in booth-to-booth connections were caused by human errors in static configurations.

continued on next page

Lessons Learned for OSI at INTEROP *(continued)*

In contrast, the show floor configurations for the Internet needed only a few entries. Since OSI routing allows many network addresses to be summarized by a shorter address string (or network prefix), the static configurations for the Internet were few (2–4) and of short length (1–2 octets).

The use of IS–IS within the show network would greatly reduce the amount of effort needed to support OSI traffic. Since INTEROP 91 Fall, NASA Science Internet (NSI) has employed OSPF for IP and IS–IS for OSI. NSI added IS–IS to its network running OSPF and encountered few problems. Using IS–IS for INTEROP 92 seems within today's technical reach.

•*Lesson 6: The Challenge of the INTEROP Show Floor has Grown.* Success in networking is the best of all times and the worst of all times. The INTEROP show floor has grown from a few machines hooked together by 2 or 3 routers on an Ethernet to a network spanning all possible technologies with over 25 backbone routers. What other networks do in months, INTEROP tries to do in 3 days. The OSI demonstration encountered the problems you would expect in such an environment. The physical and logical connections within the booth and to the outside needed to be made prior to any network testing. Little things like power coming in a few hours behind schedule become critical in this hectic atmosphere. Demonstrations which are expected to work in all three areas: within a booth, between booths, and between the booth and sites on the network need every second of testing time. Also, in an Infrastructure- or Internet-wide demo a large number of people are needed to solve a problem. A great deal of scheduling needs to take place to allow quick debugging. The network connection needs to be ready at least a full day in advance of the show to give an Infrastructure demonstration time to check out the demo.

Time constraints

The INTEROP 91 OSI demonstration uncovered a need for improvement in the show floor scheduling and debugging. Infrastructure debugging sessions were scheduled and dismissed due to the lack of network connectivity. The final OSI link to the Internet came up within an hour of show time. Fortunately, due to lots of pretesting, most things worked. But an hour is just not enough time to fix any problems.

Several excellent volunteers helped set up INTEROP 91 Fall. People who know how to set up IP networks worked long and hard on the shownet. Was it the static configurations for OSI that slowed network set up down? Was it sheer amount of physical set up? Was it lack of wide spread knowledge of OSI? Was it the last minute switch of vendor for some of the show floor routers? Was it something else? Improvement is needed for INTEROP 92.

Tools

One improvement OSI vendors need to make is a common set of network tools. Network tools include an OSI *ping,* an OSI *traceroute,* and a listing of OSI network routing tables. While we expect these functions to work across all IP hosts and routers, this functionality is not available on every OSI host. Most OSI routers provide these network tools. However, not all OSI *pings* and *traceroutes* interoperate between routers. As a result of INTEROP 91 Fall the IETF *Network layer OSI OPerational* (NOOP) working group is preparing an RFC on OSI network tools.

• *Lesson 7: TCP/IP versus OSI—Are we learning or emoting?* The friendly competition between TCP/IP and OSI has strengthened both protocol suites. The strength of TCP/IP and the Internet has been the "make it work" attitude. The ISO protocols are agreements between many nations. Both groups have something they can learn from the other protocol suites' successes and failures.

Best and worst of times

When this friendly competition and bantering gives way to unthinking emotional arguments, we all cease learning. While working on INTEROP 91 Fall, I saw people with IP backgrounds learn how to make OSI work in their networks. To my delight, people from OSI backgrounds or companies learned a lot about the Internet and IP. Sadly, I also witnessed some of the most close-minded emotional arguments about OSI and TCP/IP. I salute those IP people who took time to strengthen TCP/IP by learning about OSI. I salute those OSI people who made OSI stronger by learning about TCP/IP and the Internet. I suppose as always, this INTEROP was the best of all times and the worst of all times.

References

[1] Jacobsen, O., "INTEROP 88 Conference Report," *ConneXions,* Volume 2, No. 11, November 1988.

[2] Almquist, P., "The INTEROP 88 Network—behind the scenes," *ConneXions,* Volume 3, No 2, February 1989.

[3] Dern, D., "Highlights from INTEROP 89," *ConneXions,* Volume 3, No. 11, November 1989.

[4] Knowles, S., "The INTEROP 89 Network: From one of its builders," *ConneXions,* Volume 4, No. 2, February 1990.

[5] Dern, D., "INTEROP 90 Wrap-Up: A View from the Floor," *ConneXions,* Volume 4, No. 11, November 1991.

[6] Knowles, S., "Building the INTEROP 90 Show Network," *ConneXions,* Volume 5, No. 9, September 1991.

[7] Dern, D., "'Dear Cliff': A Report From INTEROP 91 Fall," *ConneXions,* Volume 5, No. 12, December 1991.

[8] *ConneXions,* Volume 3, No. 8, August 1989, "Special Issue: Internet Routing."

[9] *ConneXions,* Volume 5, No. 1, January 1991, "Special Issue: Inter-domain Routing."

[10] "Components of OSI," a series of articles in *ConneXions.*

SUSAN HARES has her B.S.E. in Computer Engineering from the University of Michigan. She works for Merit on the NSFNET project. Susan is active in IETF and ANSI standards committees. Susan is a network warrior on her third tour of duty. Her first tour of duty was with a private network which supported X.25, lots of terminal servers, and Bisync. The network supported the "back-end" processing of the stock market sales. People relied on it for financial information to make investments. When things failed, bank officers lost out on investments and yelled at the nearest network person. Enduring this verbal abuse, helped Susan learn the value of reliable networks. Her second tour of duty was in the MAP/TOP OSI world. She participated on standard committees writing the MAP/TOP 3.0 specification. She also helped develop products for MAP/TOP protocol suite. Susan learned a lot about what makes OSI succeed or fail from this tour of duty. The company Susan worked for made industrial computers. To encourage each engineer to do their best to make all network products (software and hardware) reliable, a true story was told to each Engineer. This story detailed how a computer error caused a computer controlled robot to put a bucket of molten metal through the wall of a factory. Susan's latest tour of duty has been with Merit on the NSFNET project. During the early days she worked on Routing Coordination for the TCP/IP Internet. Lately, Susan has been working on OSI in the NSFNET and the Internet. During INTEROP 91 Fall in October 91, she helped organize a very successful Internet-wide OSI demonstration between 30 companies. She can be reached as `skh@merit.edu`.

BREAKING THE BARRIERS

Metamail—Multimedia Mail for the Masses

On behalf of Bellcore, I am happy to announce the availability of the *metamail* software to the e-mail community. This package, which is available free of charge for unlimited use by anyone for any purpose, is offered in the hope of making multimedia mail more widespread.

Background

The basic idea of "multimedia" electronic mail is to extend e-mail as we now know it to include many other types of data beyond plain English text. In particular, there is no reason, in principle, why e-mail should not include text in any of the world's languages and character sets, nor why e-mail should not include pictures, sounds, animations, active spreadsheets, or any other kind of data that can be stored on a computer.

In recent years, various research systems and even some commercial products have extended e-mail to include some or all of these capabilities. Until recently, however, none of them worked together, and all of them required whole communities of users to abandon their old tools en masse in favor of the new tools of a single software vendor.

Recent developments have the promise of changing all of that. There is a new proposed standard for the format of multimedia mail, which would make software from different vendors able to work together smoothly with multimedia mail, as they do now with plain text mail. The software being announced here implements that proposed standard, but takes it a step further by incorporating it into the existing tools with which people read mail today, allowing multimedia mail to be adopted in an evolutionary rather than a revolutionary fashion.

MIME

Metamail is a package that can be used to convert virtually *any* mail-reading program on UNIX into a multimedia mail-reading program. It is an extremely generic implementation of MIME (*Multipurpose Internet Mail Extensions*), the proposed standard for multimedia mail formats on the Internet. The implementation is extremely flexible and extensible, using a *mailcap* file mechanism for adding support for new data formats when sent through the mail. At a heterogeneous site where many mail readers are in use, the mailcap mechanism can be used to extend them all to support new types of multimedia mail by a single addition to a mailcap file.

Viewers

The core of the package is a mechanism that allows the easy configuration of mail readers to call external "viewers" for different types of mail. However, beyond this core mechanism, the distribution includes viewers for a number of mail types defined by the MIME standard, so that it is useful immediately and without any special site-specific customization or extension. Types with built-in support in the *metamail* distribution include:

- Plain US ASCII (i.e., English) text, of course.
- Plain text in the ISO-8859-8 (Hebrew/English) character set.
- Richtext (multifont formatted text, termcap-oriented viewer).
- Image formats (using the *xloadimage* program under X11).
- Audio (initial "viewer" for SPARCstations).
- Multipart mail, combining several other types.
- Multipart/alternative mail, offering data in multiple formats.
- Encapsulated messages.
- Partial and external messages (for large data objects).
- Arbitrary (untyped) binary data.

Other media types and character sets may be easily supported with the mailcap mechanism, using the provided types as examples/templates. The *metamail* software also provides rudimentary support for the use of non-ASCII characters in certain mail headers, as described by a companion document to the proposed MIME standard.

The *metamail* distribution comes complete with a small patch for each of over a dozen popular mail reading programs, including Berkeley *mail, mh, Elm, Xmh, Xmail, Mailtool, Emacs Rmail, Emacs VM, Andrew,* and others. Crafting a patch for additional mail readers is relatively straightforward.

Retrieval and installation

In order to build the *metamail* software, a single "make" command followed by a relatively short compilation will suffice. Patching your mail reader is somewhat harder, but can usually be accomplished in less than an hour if you have the sources at hand. The experience of beta testers is that the *metamail* package can easily be used to get multimedia mail working with your existing mail readers in less than half a day.

To retrieve the file, use anonymous FTP to the machine `thumper.bellcore.com` (Internet address 128.96.41.1). Type "`cd pub/nsb`." In that directory, you will find:

- `mm.tar.Z` This is a compressed tar file containing the entire *metamail* distribution. Uncompress it, untar it, and read the top-level "README" file for further instructions. Strictly speaking, this is the only thing you really need to retrieve.
- A subdirectory called "`samples`." Except for the README file, each file in this directory is a sample MIME-format message, which can be used to test your *metamail* installation.
- `BodyFormats.{ps,txt,ex}`. A copy (in *PostScript/text/Andrew* format) of the latest draft of the MIME proposed standard. This document is also available as an Internet Draft.
- `Configuration.{ps,txt,ex}`. A copy (in *PostScript/text/Andrew* format) of the latest draft of the Internet informational RFC describing the mailcap file format. This document is also available as an Internet Draft.

Mailing list

A new mailing list has been set up for discussion of the *metamail* software and related issues. The mailing list is:

`INFO-MM@thumper.bellcore.com.`

Requests to join the list should be directed to:

`INFO-MM-REQUEST@thumper.bellcore.com.`

Please feel free to recirculate this announcement as widely as possible.

—*Nathaniel S. Borenstein*
<nsb@bellcore.com>
Member of Technical Staff, Bellcore.

[Ed.: See also the article "Multimedia Mail From the Bottom Up—or Teaching Dumb Mailers to Sing," by Nathaniel S. Borenstein, in *ConneXions,* Volume 5, No. 11, November 1991.]

RFC-Info Service now available

by Jon Postel, USC-ISI

RFC-Info is an e-mail based service to help in locating and retrival of RFCs and FYIs. Users can ask for "lists" of all RFCs and FYIs having certain attributes ("filters") such as their ID, keywords, title, author, issuing organization, and date. Once an RFC is uniquely identified (e.g., by its RFC number) it may also be retrieved.

How to use

To use the service send e-mail to `RFC-INFO@ISI.EDU` with your requests in the body of the message. Feel free to put anything in the Subject; the system ignores it. (All is case independent, obviously.) To get started you may send a message with requests such as in the following examples (without the explanation between [and]):

`Help: Help`	[to get this information]
`List: FYI`	[list the FYI notes]
`List: RFC` `Keywords: window`	[list RFCs with "window" as keyword or in title]
`List: FYI` `Keywords: window`	[list FYIs about windows]
`List: *` `Keywords: window`	[list both RFCs and FYIs about windows]
`List: RFC` `Title: ARPA*NET`	[list RFCs about ARPANET, ARPA Network, etc.]

`List: RFC`
`Organization: MITRE`
`Dated-after: Jan-01-1989`
`Dated-before: Dec-31-1991`
	[list RFCs issued by MITRE, dated 1989–1991]

`List: RFC` `Obsoletes: RFC0010`	[list RFCs obsoleting a given RFC]
`List: RFC` `Author: Bracken*`	[list RFCs by authors starting with "Bracken"] [* is a wild card matches everything]

`List: RFC`
`Authors: J. Postel`
`Authors: R. Gillman`
	[list RFCs by both Postel and Gillman]
	[note, the "filters" are ANDed]

`List: RFC` `Authors: Crocker`	[list RFCs by any Crocker]
`List: RFC` `Authors: S.D. Crocker`	[list only RFCs by S.D. Crocker]
`Retrieve: RFC` `Doc-ID: RFC0822`	[retrieve RFC-822] [note, always 4 digits in RFC#]
`Help: Manual`	[to retrieve the long user manual, 30+ pages]
`Help: List`	[how to use the LIST request]
`Help: Retrieve`	[how to use the RETRIEVE request]
`Help: Topics`	[list topics for which help is available]
`Help: Dates`	["Dates" is such a topic]
`List: keywords`	[list the keywords in use]
`List: organizations`	[list the organizations known to the system]

Please try using this service. Report problems to `RFC-MANAGER@ISI.EDU`.

Call for Participation

Toward a Truly Global Network, a Section of the 36th Annual Meeting of the International Society for the Systems Sciences will be held at the University of Denver, in Denver, Colorado, July 12–17, 1992.

Background

Computer-mediated communication networks are proliferating and growing rapidly. They are an increasingly important part of the global communication infrastructure, yet these networks are not truly global —they are concentrated in North America, Western Europe, and parts of Asia. Since networks increase industrial and intellectual productivity and efficiency, they may widen the gap between "have" and "have-not" nations.

However, there are appropriate-technology networks, for example, *Relcom* in the Soviet Union, in lesser developed nations. The purpose of this section is to bring together people operating networks like Relcom and scholars interested in truly global networking.

Topics

We will have broad participation, with papers from nuts-and-bolts to the visionary. Topics will include descriptions of networks, technology, applications, and social, economic and political considerations and implications.

ISSS

Toward a Truly Global Network, will be a multi-session section at the *International Society for the Systems Sciences* (ISSS) Annual Meeting. Other sections will address a variety of social and technical systems.

More information

Further information on this section can be obtained from:

Professor Larry Press
CSUDH
10726 Esther Avenue
Los Angeles, CA 90064
USA
Phone: +1 310-475-6515
Fax: +1 310-516-3664
Internet: `lpress@venera.isi.edu`

or

Professor Kjell Samuelson
Royal Institute of Technology
Narvegen 8
Stockholm 11523
SWEDEN
Phone: +46 8-663-6302

Further information on ISSS and the meeting in general contact:

Professor J. D. R. de Raadt, Managing Director
ISSS, College of Business
Box 8793
Idaho State University
Pocatello, ID 83209
USA
Phone: +1 208-233-6521
Fax: +1 208-236-4367

Call for Papers

NETWORKS 92—An International Conference on Computer Networks, Architecture and Applications will be held 26–28 October, 1992 in Trivandrum, India. The conference is jointly sponsored by the International Federation for Information Processing (IFIP) and the Computer Society of India.

Topics

The conference provides an international forum for presentation of recent results in the general area of Networks with special emphasis on Applications, issues relating to Management, Security and Performance. More specifically, the topics of interest include, but are not limited to:

Protocol Related Issues:

- Multimedia Applications
- Distributed Applications
- Distributed Information Management
- Client–Server Models
- Task Allocation and Load Balancing
- Specification and verification of protocols

Performance Related Issues:

- Quality of Service (QOS) Measurement
- Management of QOS
- Higher Layer Protocol related performance Issues
- Workload Characterization of Multi Media Applications

Technology Related Issues:

- Implementation Issues relating to Multi Media Applications
- High Speed Networks Implementations
- Wireless Network Implementation
- Relevance to Developing Countries

Format

The conference is single-track. Only papers of exceptional quality will be accepted for presentation at the conference and inclusion in the Proceedings. There may be a few invited papers of immediate relevance to the conference. There will be four tutorials on current topics on October 25th, 1992.

Information for authors

Prospective Authors are invited to submit full papers concerned with both Theory and Practice. The areas of interest for the conference include, but are not limited to topics mentioned above. Authors should note the following:

- All submissions are refereed.
- Papers should be double spaced, should be less than 20 pages and should have an abstract.
- There should be a cover page giving Title, Author(s), Complete Address and Affiliation, Telephone Numbers, Fax numbers and Electronic Mail Addresses.

- Authors of accepted papers will be expected to sign a copyright release form.
- Participants copy of the proceedings will be preprinted and distributed during the conference.
- Submit *five* copies of each paper to the Program Chair in your region.
- Send in your "intention to submit a paper" information (Title, Author(s), Address, Telephone, Fax, E-mail address) to:

  `NET92@shiva.ernet.in`

Regional Program Chairs

North America:

G.V. Bochmann
Universite De Montreal
Department D'Informatique et Recherche Operationnelle
Case Postale 6128
Succursale A
Montreal
Quebec H3C 3J7
CANADA
Tel.: +1 514-343-2155
Fax: +1 514-343-2155
E-mail: `bochmann@.iro.Umontreal.CA`

Europe:

Guy Pujolle
Professor
Department of Computer Science
Universite Pierre et Marie Curie
Laboratoire MASI
4 Place Jussieu
Paris Cedex 05,75252
FRANCE
Tel.: +33 1 39 02 01 09
Fax: +33 1 30 21 19 83
E-mail: `gp@masi.ibp.fr`

Asia / Africa / Australia:

S.V. Raghavan
Department of Computer Science & Engineering
Indian Institute of Technology
Madras
INDIA
Tel.: + 91 44 235-2377 and 235-0563
Fax: + 91 44 235-0509
E-mail: `raghavan@shiva.ernet.in`

Important dates

Deadlines for Paper Submission:	April 30, 1992
Notification of Acceptance:	July 15, 1992
Camera Ready Paper Due:	September 15, 1992

Call for Participation and Pre-Announcement

The *OpenForum'92* "Distributed Computing: Practice and Experience" technical conference will be held November 25–27, 1992 at the Royal Dutch Fairgrounds, in Utrecht, The Netherlands. *OpenForum'92* focuses on distributed computing and features keynote addresses, technical presentations, panel discussions and a work-in-progress session. The opening keynote address will be given by Professor Andrew S. Tanenbaum from Vrije Universiteit, Amsterdam.

Topics

Original contributions are sought that address fundamental issues of distributed computing. Topics of interest include, but are not limited to:

- Distributed system management
- Reliability and availability in distributed systems
- Security in open distributed environments
- Microkernel experiences
- Operating system support for multimedia
- Languages for distributed computing
- Communication issues of distributed computing
- Distributed file systems
- Distributed applications

Submissions

Please submit an extended abstract, not exceeding 3000 words, by post or e-mail to the Chair of the technical conference programme by March 15th, 1992. Each submission must provide sufficient detail to allow the program committee to assess the merits of the contribution. Submitted abstracts must indicate clearly their relevance and contribution to the overall theme of the conference.

Work-in-progress sessions

Authors are invited to submit work-in-progress (WIP) proposals limited to 3000 words or 5 pages. Please submit WIP proposals to the Program Committee Chair by post or e-mail. These sessions provide speakers with 10 minutes to speak on current work and receive valuable feedback. All submissions should be sent to:

OpenForum'92 Technical Conference Chair
Dag Johansen
Dept. of Computer Science
University of Tromsø
N-9000 Tromsø
NORWAY
Tel.: +47 8 34 40 47
Fax: +47 8 34 45 80
E-mail: `dag@cs.uit.no`

Important dates

Extended abstract submissions due:	March 15, 1992
Acceptance notification:	May 10, 1992
Deadline for final papers:	August 15, 1992
Work-in-Progress submissions due:	November 1, 1992

Programme Committee

Dag Johansen, University of Tromsø, Norway (Chair)
Lori Grob, Chorus systèmes, France
Heinz Lycklama, Interactive Systems, USA
Robbert van Renesse, Cornell University, USA
Volker Tschammer, GMD-FOKUS, Germany

Announcement and Call for Papers

The *Third USENIX UNIX Security Symposium* will be held September 14–16, 1992 in Baltimore, Maryland. The symposium is being hosted by USENIX In cooperation with The Computer Emergency Response Team (CERT).

Goal

The goal of this symposium is to bring together security practitioners, system administrators and system programmers, and anyone with an interest in computer security as it relates to networks and the UNIX operating system. The symposium will consist of tutorials, invited speakers, technical presentations, and panel sessions.

Format

This will be a three-day, single-track symposium. The first day will be devoted to tutorial presentations. The following two days will include technical presentations and panel sessions. There will also be two evenings available for birds-of-a-feather sessions and work-in-progress sessions.

Topics

Papers are being solicited in areas including but not limited to:

- User/system authentication
- File system security
- Network security
- Security and system management
- Security-enhanced versions of the UNIX operating system
- Security tools
- Network intrusions (including case studies and intrusion detection efforts)

Important dates

Extended abstracts due:	May 15, 1992
Program Committee decisions made:	June 15, 1992
Camera-ready papers due:	July 31, 1992

Paper submissions

Send seven copies of each submission to the program chair:

Edward DeHart
Computer Emergency Response Team
Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213-3890
+1 412-268-6179
`ecd@cert.sei.cmu.edu`

Call 1-800-INTEROP or 1-415-941-3399 today for more information or to reserve your space.

INTEROP® 92 SPRING

18–22 May 1992 • Washington, D.C. Convention Center

Call for Papers

The IEEE *Journal on Selected Areas in Communications* announces a forthcoming issue on *Network Management and Control.* The issue will focus on modeling, structure, and representation of management information and the operations and control aspects of network control and management.

Further, papers focusing on issues concerning the design and implementation of the management information base, the convergence towards common management tools and environments and examples of managing specific technologies and/or functions are also solicited. Both tutorial and original contributions will be accepted.

Topics

Topics of interest include the following:

- Management Information Modeling, Structure and Representation
- Knowledge Representation and Reasoning in Management
- Management Standards and Architectures
- Formal Description of Management
- The Integration of Real-Time Control and Management
- Management of specific technologies such as Telecommunication Networks, Distributed Systems, and Data Communications Networks
- Management Functions:
  Configuration, Fault, Performance, Accounting and Security

Important dates

Prospective authors are requested to submit five (5) copies of their manuscript to one of the guest editors listed below. Deadline of the initial paper submission is June 1st, 1992. Acceptance notification by December 1, 1992. Final papers are due February 1, 1993, with an anticipated publication date in the third quarter of 1993.

Guest Editors

Aurel A. Lazar
Columbia University
Dept. of Electrical Engineering and
Center for Telecommunications Research
New York, NY 10027-6699
USA
Tel: +1 212-854-1747
Fax: +1 212-932-9421
E-mail: aurel@ctr.columbia.edu

Roberta Cohen
1L333 AT&T Bell Labs
101 Crawfords Corner Rd
P.O. Box 3030
Holmdel, NJ 07733-3030
USA
+1 908-949-2203
+1 908-949-8569
rsc@arch3.att.com

Chris Sluman
Sema Group
Regal House
14 James Street
London WC2E 8BT
ENGLAND
Tel: +44 71-497-0519
Fax: +44 71-497-2155
E-mail: cs2@hermes.mod.uk

Mark Sylor,
Digital Equipment Corp.
110 Spitbrook Road
ZK01-3/J35
Nashua, NH 03062-2642
USA
+1 603-881-2692
+1 603-881-0120
sylor@blumon.enet.dec.com

Book Review

Communications for Cooperating Systems: OSI, SNA, and TCP/IP, by R. J. Cypser, Addison-Wesley's Systems Programming Series, 1991, ISBN 0-201-50775-7, 743 pages, cloth.

IBM watchers should already be familiar with Dr. Cypser from his earlier book "Communications Architecture for Distributed Systems." The former book was published in 1978 and served as an exemplary introduction to the context and philosophy which guided IBM's initial SNA design. Even though SNA has significantly evolved from its 1978 offering, Dr. Cypser's earlier book still remains relevant and informative.

IBM's direction

Dr. Cypser's newest book has continued in the earlier book's tradition by describing IBM's current directions in data communications. This newest book describes an IBM viewpoint of a multivendor world in which SNA, OSI, TCP/IP, and NetBIOS must co-exist. While the book does not claim to present an official IBM position of such a universe, it does reflect an IBM viewpoint gathered from Dr. Cypser's long tenure within IBM. Previous to his retirement in 1989, Dr. Cypser was director of technical communications for IBM. That experience, together with extensive conversations with hundreds of IBM technical staff members located in virtually every IBM development laboratory world-wide, have formed the basis for the viewpoint presented in this book. Should further evidence of the author's IBM credentials be needed, one need only mention that the book was written at Ellen Hancock's (IBM's Vice President and General Manager of Communications Systems) encouragement and was actively supported by both John Hunter (IBM's Director of Communications Systems Architecture and Technology) and Rick McGee (IBM's Manager of Communications Architecture).

OSNA

The book's thesis is an architected multi-protocol environment named the *Open Systems Network Architecture* (OSNA). OSNA is essentially an IBM *Systems Application Architecture* (SAA)-based "macro architecture for heterogeneous systems with an evolutionary thrust toward multiprotocol interoperability and standards." End systems residing in this environment belong to one of four node types:

- *Pure SNA nodes:* Pure SNA nodes include both traditional SNA nodes (referred to as "Subarea SNA") as well as *Advanced Program to Program Networking* (APPN) nodes. Several chapters are devoted to describing both of these environments. The clear picture which emerges is that APPN will continue to grow in importance with new APPN product offerings. Once VTAM and NCP have been enhanced with APPN capabilities, the two environments will be effectively integrated. (This is anticipated to occur in 1992.) Once this happens, an APPN node's Control Point and Subarea SNA's SSCP will essentially become peers of each other. The net result should be a metamorphosis of SNA with APPN characteristics such as dynamic reconfigurability, reduction (or elimination) of the need for Sysgens, and non-disruption of the network during network changes.

- *Combined OSI/SNA nodes*: The goal of the second node type is to fully function in either an SNA environment or an OSI environment. "IBM's integration strategy is to implement native, fully compliant OSI products while maximizing interoperability, resource sharing, and the value-added functions of SNA."

continued on next page

Book Review *(continued)*

This is accomplished through extending the commonality of features between SNA and OSI in principally four areas:

1: Both would support the SAA *Common Programming Interface* (CPI) Application Programming Interfaces (APIs). In the data communications domain, this would mean common support for the IBM CPI-C API.

2: Both would support common application services.

3: Both would possess a common directory service.

4: Both would support the same common transport providers. These transport providers are connection-oriented either at the data link layer (e.g., IEEE 802.2 *Logical Link Control Type 2* (LLC2) connectivity) or at the network layer (e.g., X.25). Note, however, that no OSI Profile in the United States (e.g., US GOSIP) currently permits protocol stacks of the former composition.

OSI/CS

The product family which is based upon the IBM *OSI/Communications Subsystem* (OSI/CS) product is targeted to evolve into a combined OSI/SNA node. That is, the IBM application services which are designed to use OSI/CS (e.g., IBM's *OSI/File Service* product and IBM's various CCITT X.400-related products [ONDS, XPC, XDC]) are able to function as gateways linking the SNA world to the OSI world. Similarly, OSI applications in two different OSI networks are able to communicate with each other over an intermediate SNA network through the services provided by OSI/CS. While this family is US GOSIP-conformant, its decidedly preferential orientation is for connection-oriented data link or network layer communications. Finally, the OSI/CS product has a CCITT X.500-based directory capability which is able to address both SNA and OSI nodes.

- *Pure OSI nodes:* These are envisioned primarily to be non-IBM OSI nodes. These nodes may function in a pure OSI environment or be linked to the SNA world via the combined OSI/SNA nodes.
- *Other combined nodes:* Two protocols are explicitly mentioned in the "other" category: *TCP/IP* and *NetBIOS*. References to NetBIOS were explicitly stated to solely mean the "IBM NetBIOS" whose stack is the NetBIOS protocol directly located over the data link layer—as opposed to the RFC 1001/1002-conformant version of NetBIOS which has been blessed by X/Open. Also, although many TCP/IP protocols (e.g., IP, TCP, ARP, SNMP) are referenced for explanatory comparisons, TCP/IP itself is usually referenced in the context of the Open Software Foundation's (OSF) *Distributed Communications Environment* (DCE).

Integration

The book is primarily concerned with SNA and the Integrated OSI/SNA approach. It is these environments upon which the primary multiprotocol integration occurs and around which OSNA is architected. While the other two node types receive comparatively little attention, they are by no means ignored. Rather, they are to be integrated into OSNA whenever possible via four different mechanisms:

1: *Gateways:* The dominant method of interconnecting diverse environments into OSNA is by means of application-layer gateways. So many gateways are enumerated in this context that it is not possible to mention all of them here.

2: *Common Application Programming Interfaces* (APIs): IBM's CPI-C API and OSF's Remote Procedure Call (RPC) were explicitly mentioned in this context.

3: *Common Managers:* The PS/2 Communications Manager was mentioned as a platform upon which divergent protocol environments can be integrated.

4: *Multiple Protocol Passthrus together with protocol encapsulation or conversion:* This would permit "foreign" protocol stacks to traverse SNA networks.

A different world-view

Perhaps it goes without saying that the world-view reflected by this book is substantially different from the world-view of the average TCP/IP user. This can be seen by the limited emphasis given to TCP/IP and the SNA-centric orientation. More telling, however, are the explanatory sections which introduce basic data communications technologies. For example, Dr. Cypser argues forcibly for the advantages associated with connection-oriented network services—and the disadvantages of connectionless network services (such as IP). Dynamic routing, which is presumed by both TCP/IP and US GOSIP, is unfavorably described.

SNA bias

Similarly, when introducing directory services, a strong SNA bias to the introductory material is seen: "In all cases, the [ascertainment of the] route is the desired result [of a directory lookup]." The text also explains that up to three layers of directory functionality (i.e., network-layer and data link-layer directories, in addition to the traditional application layer directory service) are needed to assist in the route selection process.

These biases shed considerable light concerning the data communications orientation of many influential people within IBM. In addition, they indicate the probable direction future IBM data communication product offerings will take unless they are otherwise guided by outside market forces.

—Eric Fleischman

Write to *ConneXions*!

Have a question about your subscription? Are you moving, and need to give us your new address? Suggestions for topics? Want to write an article? A letter to the Editor? Have a question for an author? Want to enquire about back issues? (there are now over sixty to choose from; ask for our free 1987–1991 index sheets). We want to hear from you. Simply write, call, fax, or e-mail us at:

ConneXions—The Interoperability Report
480 San Antonio Road, Suite 100
Mountain View, CA 94040–1219
USA
Phone: +1 415-941-3399 or 1-800-INTEROP (Toll-free in the USA)
Fax: +1 415-949-1779
E-mail: `connexions@interop.com`

Next month:

A *Special Issue* on Emerging Broadband Services.

CONNEXIONS
480 San Antonio Road
Suite 100
Mountain View, CA 94040
415-941-3399
FAX: 415-949-1779

ADDRESS CORRECTION
REQUESTED

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

CONNEXIONS

EDITOR and PUBLISHER — Ole J. Jacobsen

EDITORIAL ADVISORY BOARD

Dr. Vinton G. Cerf, Vice President,
Corporation for National Research Initiatives

A. Lyman Chapin, Chief Network Architect,
BBN Communications

Dr. David D. Clark, Senior Research Scientist,
Massachusetts Institute of Technology

Dr. David L. Mills, Professor,
University of Delaware

Dr. Jonathan B. Postel, Communications Division Director,
University of Southern California, Information Sciences Institute

CONNEXIONS

Subscribe to CONNEXIONS

U.S./Canada ❑ $150. for 12 issues/year ❑ $270. for 24 issues/two years ❑ $360. for 36 issues/three years

International $ 50. additional **per year** **(Please apply to all of the above.)**

Name ________________ Title ________________

Company ________________

Address ________________

City ________________ State ________ Zip ________

Country ________________ Telephone (   ) ________________

❑ Check enclosed (in U.S. dollars made payable to CONNEXIONS).
❑ Visa ❑ MasterCard ❑ American Express ❑ Diners Club Card # ________________ Exp.Date______

Signature ________________

Please return this application with payment to:

CONNEXIONS
480 San Antonio Road, Suite 100
Mountain View, CA 94040 U.S.A.
415-941-3399 FAX: 415-949-1779

Back issues available upon request $15./each
Volume discounts available upon request